

Extensive Air Showers: Model Dependence and the Longitudinal Profile

Paul T. Mikulski

A dissertation submitted to the Johns Hopkins University in conformity with the
requirements for the degree of Doctor of Philosophy.

Baltimore, Maryland

2000

Abstract

The physics of atmospheric cascades induced by ultrahigh energy cosmic rays above the GZK cutoff is investigated. A flexible air shower generator is developed for the exploration of air shower properties under various assumptions about the underlying hadronic physics. Included is an efficient algorithm for the simulation of electromagnetic cascades which emphasizes consistency with the standard Greisen formula for photons. The uncertainties in hadronic interaction physics relevant to air shower properties are bracketed with conservative models which characterize limiting scenarios. The effects of scaling violations in hadron-nucleon and hadron-nucleus interactions are addressed. Emphasis is on the fragmentation region which is of most relevance to air shower development. The physics uncertainties imply an uncertainty in the average depth at maximum for proton induced cascades at the highest energies which is comparable to the separation between protons and iron. Consequently, it will be difficult to extract information about the composition at the highest energies. The possibility of detecting new physics or exotic primaries is discussed. Given the uncertainties associated the highest energies, very long interaction lengths are required in the initial stages of anomalous showers to clearly distinguish them from traditional scenarios (smooth evolution of the hadronic physics and a composition of protons and/or nuclei). An implementation of the LPM effect which suppresses electromagnetic cross sections in air at the highest energies is also described.

Acknowledgements

I am grateful to my advisors Gabor Domokos and Susan Kovesi-Domokos for their guidance over the full course of my graduate studies. They have always been caring and understanding not only with respect to my research but also to my life as a whole. This helped me greatly through some difficult times. I am particularly grateful to them for their guidance leading up to and during my leave of absence to teach at St. Mary's College.

I thank John Krizmanic for his support over the last two years. He took a vested interest in the nuts and bolts of my research. My discussions with him helped me pull my work into sharper focus. I would also like to thank my fellow graduate student, Dobromir Velkov, who helped me with many programming difficulties.

Many thanks to Janet Krupsaw for always being there for all the graduate students. Not only does she always get the job done without a hitch, she is one of the nicest and most caring people I know.

On behalf of Enzo, I would like to thank the members of the sixth floor and all those who made his stay at Hopkins most enjoyable.

Contents

List of Figures	vi
List of Tables	x
1 Introduction	1
1.1 The big picture	2
1.2 Overview of cosmic rays	3
1.2.1 Energies up to a TeV	4
1.2.2 The knee and the ankle	5
1.2.3 The highest energy cosmic rays	6
2 Electromagnetic Cascades	8
2.1 Basic Description	9
2.2 The Track Length Integral	11
2.3 Analytic Cascade Theory	13
2.4 The Greisen Formula	16
2.5 A Modified Greisen Formula	19
2.6 Procedure for developing parameterizations	21
2.7 Neutral pions	22
2.8 Deconvoluted photons	23
2.9 Electrons	24
2.10 Bremsstrahlung and an effective model	24
2.10.1 A simple effective splitting model	25
2.10.2 A more general effective splitting model	26
2.10.3 Choosing a model	27
2.11 Fluctuations in shower development	29
2.12 Summary	32
3 Hadronic Interactions	34
3.1 Cross sections	36
3.2 Hadronization: general considerations	41
3.2.1 The Hillas splitting algorithm	42

3.2.2	Analytic description of fragments	43
3.2.3	String fragmentation	45
3.3	The leading particle	47
3.4	The distribution of secondaries	49
3.5	Scaling violations	52
3.6	Nuclear target effects	58
3.6.1	Sampling the number of wounded nucleons	59
3.6.2	Hadronization effects	60
3.7	Pion projectiles	63
3.8	Nuclear projectiles	64
3.8.1	The wounded nucleon picture	65
3.8.2	Nuclear cross sections	66
3.8.3	The grouping of spectators	66
3.8.4	Some simulation results	67
3.9	Summary	68
4	Bootstrapping Models	70
4.1	The atmosphere	70
4.2	Basic methodology	72
4.3	The Gaisser-Hillas profile function	73
4.4	Starting depths and zenith angle	75
4.5	Some technical details	76
4.6	The fractional energy parameters	78
4.7	The depth at maximum and elongation rate	78
4.8	Bootstrapping errors	82
4.8.1	Random error	83
4.8.2	Zenith angle	84
5	Model Dependence: results and implications	86
5.1	Comparing hadronization models	87
5.2	Comparing extrapolations of cross sections	87
5.2.1	Constraining the p -air inelastic cross section	88
5.2.2	Bootstrapping results	89
5.3	Nuclear target effects revisited	90
5.4	Searching for new physics	91
5.4.1	Muons in proton showers	92
5.4.2	Prompt muons in anomalous showers	93
5.5	The LPM effect	95
5.6	Fluctuations and primary type	98
5.7	Summary	99
	Bibliography	108

List of Figures

2.1	The depth at maximum as a function of primary photon energy as prescribed by standard electromagnetic cascade theory.	16
2.2	Shower size versus atmospheric depth as described by the standard Greisen formula for a 10^{20} eV photon initiated shower.	18
2.3	Comparison of the longitudinal profiles for primaries at 10^{18} eV using an effective splitting model with $\lambda_{\text{eff}} = 0.5$. The Monte Carlo averaged 1000 showers using a simulation threshold of 10^{15} eV.	28
2.4	Results of the random search for the best effective splitting model. There is a clear minimum near $\lambda_{\text{eff}} = 0.31$	29
2.5	Comparison of the longitudinal profiles for primaries at 10^{18} eV using the best effective splitting model. The Monte Carlo averaged 1000 showers using a simulation threshold of 10^{15} eV.	30
2.6	Depth at maximum distributions for 10^{18} eV photon initiated showers. In the top histogram electrons are described by their average relative to their depth of creation. The bottom histogram uses the effective splitting model with a simulation threshold of 10^{15} eV. Each set is comprised of 1000 showers.	31
2.7	Longitudinal profiles for ten 10^{18} eV photon showers. The effective splitting model was used to conduct Monte Carlo simulation down below 10^{15} eV.	32
3.1	Energy dependence of the inelastic cross sections for hadrons on a nucleon target. The ratio of the pion to nucleon cross sections is held at a constant fixed by low energy data.	38
3.2	Energy dependence of the inelastic cross sections for hadrons on air. Because of the shadowing effects associated with the target nucleus, the ratio between the pion and nucleon cross sections is larger than for on a nucleon target.	39
3.3	Energy dependence of the interaction length for hadrons on air.	41

3.4	Leading particle distributions for nucleon projectiles on a nucleon target (top) and pion projectiles on a nucleon target (bottom). Each set is comprised of 10,000 simulated interactions.	48
3.5	x_{lab} distribution of charged pions for nucleon projectiles on a nucleon target. The histogram is constructed from 10,000 simulated interactions.	50
3.6	Two fits to the average charged multiplicity. The fit which grows logarithmically agrees well with bubble chamber data which extends up to about 5×10^{11} eV in the lab. The fit which grows as the log squared takes into account collider data which extends up to energies corresponding to about 4×10^{14} eV in the lab.	55
3.7	The average number of presplittings for the description of pp interactions in the limiting model where the forward fragmentation region is strongly correlated with the central region.	56
3.8	Energy dependence of the inelasticity in pp interactions in the limiting model where the forward fragmentation region is strongly correlated with the central region.	57
3.9	Energy dependence of the average number of wounded target nucleons, those that participate inelastically, for hadrons on air.	58
3.10	The fractional energy for presplits as a function of the number of wounded nucleons. The value at $N_w = 1$ gives the inelasticity. The lines are drawn for visual clarity, only discrete values for the number of wounded nucleons is meaningful.	61
3.11	The energy dependence of the inelasticity for protons on air. The upper curve includes scaling violations in the underlying pp interactions while the lower curve does not.	62
3.12	x_{lab} distribution of charged pions for charged pion projectiles on a nucleon target. The histogram is constructed from 10,000 simulated interactions.	63
3.13	The distribution of the number of participating nucleons from the projectile for iron at 1 TeV on air. The simulation set is comprised of 10,000 interactions.	67
3.14	The distribution of the average depth of first interaction. Each simulation event breaks up an iron nucleus of total energy 1 TeV into nucleons each at their depth of first interaction. The simulation set is comprised of 10,000 such events.	68
4.1	The energy dependence of the fractional energy parameters for nucleon induced subshowers in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. The figure shows the energy deposited into the electromagnetic cascade, the energy lost in muons and neutrinos and the energy ignored due to error in the simulation algorithm.	77

4.2	The energy dependence of the depth at maximum for nucleon and nucleus induced subshowers in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. The upper solid line corresponds to nucleons and the dotted lines correspond to nuclei with $A = 2, 4, 8, 16, 32$ and 64 . The lower solid line is for iron ($A = 56$).	79
4.3	The energy dependence of the elongation rate for nucleon induced subshowers in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections.	80
4.4	The energy dependence of the elongation rate for fragments in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections.	81
4.5	The depth at maximum and the fractional deposited energy for nucleons at 3×10^{20} eV in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. . . .	83
4.6	The energy dependence of the fractional energy lost in muons and neutrinos for nucleons in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. Starting from the topmost line, $\sec \theta = 0.5, 0.4, 0.3, 0.2, 0.1$ and 0 . . .	84
4.7	The energy dependence of the depth at maximum for nucleons in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. Depicted are six bootstrapping runs each conducted at a different zenith angle: $\sec \theta = 0.0, 0.1, 0.2, 0.3, 0.4$ and 0.5	85
5.1	The energy dependence of the inelastic pp cross section over the set of models investigated. The dotted line shows the default.	88
5.2	The energy dependence of the number of muons for thresholds of 10 GeV and 1 TeV evaluated for the bracketing models. Notice that the muon number is not strongly sensitive to the choice of cross sections .	93
5.3	The total number of particles above a specified threshold generated in anomalous interactions above the unification energy as a function of the unification energy.	94
5.4	The average profile over 1000 showers for protons at 10^{20} eV. The starting depth and development phase are averaged separately so the resulting shape is more characteristic of individual showers. The hadronization model uses default cross sections with scaling violations in pp interactions and nuclear target effects turned on.	96

5.5	The depth at maximum distribution for protons and iron at 10^{20} eV over 1000 showers in each set. Iron is not sensitive to the LPM effect so only one histogram is shown in this case. The hadronization model uses default cross sections with scaling violations in pp interactions and nuclear target effects turned on.	97
5.6	Using default cross sections, the energy dependence of the depth at maximum and the elongation rate for nucleons with only nuclear target effects turned on. The dashed lines in the depth of maximum plot are for iron. The results are compared with the baseline model, the upper line of each pair.	101
5.7	Using default cross sections, the energy dependence of the depth at maximum and the elongation rate for nucleons with only scaling violations in pp interactions turned on. The dashed lines in the depth of maximum plot are for iron. The results are compared with the baseline model (the upper line of each pair).	102
5.8	Using default cross sections, the energy dependence of the depth at maximum and the elongation rate for nucleons with scaling violations in pp interactions and nuclear target effects turned on. The dashed lines in the depth of maximum plot are for iron. The results are compared with the baseline model, the upper line of each pair.	103
5.9	The energy dependence of the inelastic cross section (top) and the interaction length (bottom) for nucleons on air over the set of models investigated. The dotted line shows the result for default cross sections.	104
5.10	The energy dependence of the depth at maximum for the scaling model of hadronization (top graph) and with scaling violations in pp interactions and nuclear target effects turned on (bottom graph) over the set of models investigated. The dotted line shows the result for the default cross sections.	105
5.11	The energy dependence of the depth at maximum distribution with only nuclear target effects turned (top graph) and with only scaling violations in pp interactions turned on (bottom graph) over the set of models investigated. The dotted line shows the result for the default cross sections.	106
5.12	Bracketing the uncertainties in air shower development assuming proton primaries, energy dependence of the depth at maximum and the elongation rate. The dashed line is a lower bound from theory, the lower solid line is the lower bound constraint from air shower data.	107

List of Tables

2.1	Modified Greisen parameters and the average fractional energy of the typically lower energy particle for effective electromagnetic splitting models characterized by an effective mean free path. The error represents the discrepancy in area overlap between the simulated and the target profiles.	27
2.2	Summary of the modified Greisen parameters for a variety of initial conditions. There is no error for the first entry since it is an assumption by which the remaining are derived.	33
3.1	Average multiplicity of particles observed in the forward fragmentation region as observed by the LEBC-EHS collaboration for pp interactions at 400 GeV in the lab frame.	51
4.1	Summary of pion properties needed for the calculation of decay depths in the atmosphere.	72

Chapter 1

Introduction

Future cosmic ray experiments (Auger [1], OWL [2]) will definitively settle the question of the cosmic ray energy spectrum above 10^{20} eV. The fact that a handful of such events has been observed (Fly's Eye [3], AGASA [4], Haverah Park [5] and Yakutsk [6]) has presented the most puzzling problem in the field of cosmic ray physics: theory predicts they should not be there! With the increased statistics obtainable by these future experiments, it will be possible to determine the shape of the energy spectrum and any patterns in arrival directions and possibly to infer information about the nature of the composition as well.

At these energies the primary cosmic rays are observed indirectly through the shower of secondary particles they generate in the atmosphere. While a determination of energy and arrival direction of the primary is nearly model independent, a detailed understanding of the development of these extensive air showers will be required if there is to be any hope of inferring information about composition. The focus of this work is to acquire a clear understanding of how these cosmic ray induced showers develop. The primary goal is to build and utilize a framework for investigating model dependence associated with the uncertainties in the relevant physics. Two important questions with respect to future air shower observations above 10^{20} eV are addressed:

1. If a high statistics observation is obtained, will it be possible to do a meaningful analysis of the mass composition?

2. What constitutes an unusual signal that clearly indicates new physics or an exotic primary composition (something different from the protons and nuclei which dominate the cosmic ray spectrum at lower energies)?

The main tool which is systematically developed is a flexible air shower generator which can be easily modified to simulate air shower development under varying assumptions about the underlying physics. For those wishing to make use of this generator, understanding this work is the first place to begin.

The remainder of this chapter breaks into three sections. The first gives a brief summary of the experimental and theoretical agenda concerning future work on the highest energy cosmic rays. This includes a brief discussion of the phenomenon that makes this regime so puzzling: the GZK effect. The next section is an overview of what is presently known about cosmic rays extending up to but not including energies beyond 10^{20} eV. Following this, a separate section is devoted to a brief consideration of the highest energy events observed so far.

1.1 The big picture

Future experiments hope to form a clear picture of cosmic rays above 10^{20} eV which arrive at the top of the earth's atmosphere. Measurements will include a determination of the energy spectrum, an understanding of any pattern in arrival directions and hopefully information about the composition of the cosmic rays (are they protons, nuclei, something exotic?). Convolved with information about composition is information about hadronic interactions at these energies: in the collision of a primary cosmic ray nucleon and an air-nucleus the center of mass energy is over 400 TeV, more than two hundred times the energy available to collider experiments.

With a high statistics observation, it is hoped that a theory can be constructed which is consistent with observation (energies, arrival directions, composition). This theory must address not only the question of the acceleration mechanisms by which cosmic rays achieve these enormous energies, but also explain the propagation of cosmic rays from their sources to the earth. At present even a plausible picture which

can address their existence is elusive. The sources which seem capable of the required acceleration lie too far away. Interactions with the cosmic microwave background radiation, a thermal bath of photons at 2.7 K which permeates all of space, degrade the energies of protons above about 5×10^{19} eV as they propagate through space. This requires sources to be within about 100 Mpc if protons are to have a reasonable chance of making the journey to Earth without a severe degradation of their energies. This is a cosmologically near scale within which there are no known sources capable of generating these highest energies. This energy threshold of about 5×10^{19} eV, known as the Greisen-Zatsepin-Kuzmin or GZK cutoff was first pointed out by Greisen [7] and independently by Zatespin and Kuzmin [8] in the 60's. It has been carefully treated in a number of recent detailed simulations [9, 10].

If there are no cosmologically near sources, the energy spectrum should continue only with greatly reduced intensity beyond the GZK cutoff. Simply put, the center of mass energy of an ultra-high energy proton with an energy at the GZK cutoff and a typical photon from the microwave background radiation sits near the Δ resonance, an excited state of the proton which decays producing a nucleon and a pion. After an interaction, the resulting nucleon's energy lies below the GZK cutoff; it can propagate freely since kinematically the production of a pion is not allowed. At energies sufficiently above the GZK cutoff, interactions can produce multiple pions. Since the center of mass energies in these interactions are rather modest, this effect has been studied in detail in the laboratory. Calculations of the cosmic ray spectrum for various source distributions may differ slightly in detail; however for energies greater than 10^{20} eV, there does not seem to be any way around the conclusion that if the highest energy cosmic rays are protons, they must have been created within tens of megaparsecs of the earth. The effect is even more dramatic for nuclei: interactions with the same microwave background radiation break nuclei up at comparable energies [11].

1.2 Overview of cosmic rays

There are a number of very good treatments on the general subject of cosmic rays. The best reference which spans all energies is the admirable two-volume work

of Longair [12]. Works that specifically address higher energies include Sokolsky [13] and Gaisser [14]. An excellent overview with an eye on future experiments is found in the Pierre Auger Design Report [1]. The very brief treatment below is derived from these works where much greater detail can be found. In particular, the work of Sokolsky gives an excellent discussion of experimental techniques relevant to the highest energy cosmic rays.

The differential energy spectrum of cosmic rays is described well by a power law of the form

$$N(E)dE = KE^{-x}dE. \quad (1.1)$$

This power law form holds over many orders of magnitude beginning at about a GeV (10^9eV) per nucleon and extending beyond $5 \times 10^{19}\text{eV}$. The spectral index x varies somewhat as a function of energy and composition but is always in the range between about 2.5 and 3.1. Over this energy range the flux drops severely: the integrated flux of cosmic rays arriving at the top of the Earth's atmosphere above about 10^{11}eV is around one particle per square meter per second, whereas the integrated flux above 10^{20}eV is of the order of one particle per square kilometer per century.

The presence of a power law spectrum restricts acceleration mechanisms to those that can produce this shape. A model known as *Fermi acceleration* which is capable of explaining the power law form at modest energies (up to at least a PeV) involves the acceleration of particles through repeated encounters with moving magnetized plasmas. Fermi showed that on average the energy change in a particle encountering a moving plasma cloud is positive and proportional to the particle's energy [15]. After k encounters, the particle's energy will have grown to about $E_0(1+\alpha)^k$ where α is the average fractional gain in energy per encounter. This mechanism implies a power law spectrum where the spectral index is given by $x = P_{\text{esc}}/\alpha$ where P_{esc} is the probability per encounter that the particle escapes the region where acceleration can take place.

1.2.1 Energies up to a TeV

At the low end, where E refers to the kinetic energy per nucleon, up to about a TeV the primary cosmic rays can be detected directly and their energy and charge

clearly discriminated. The spectral index x takes on values in the range 2.5—2.7 for various atomic masses. Roughly 87% are protons, 12% are helium nuclei and about 1% are heavier nuclei.

The differences between species can be understood in terms of source production and propagation from sources to the earth. The relative abundances of elements produced at sources within the galaxy are modified as nuclei propagate through the interstellar medium. During their journey from source to earth they suffer collisions with the nuclei of atoms of the interstellar medium which break the nuclei into smaller fragments. This process is called *spallation*.

A fairly clear picture of the journey cosmic rays undertake in their propagation from sources to the earth can be built up by comparing the cosmic ray abundances to local solar abundances under the assumption that local solar abundances are characteristic of the abundances at sources. The comparison of stable elements gives an estimate of the amount of matter traversed. The presence of unstable elements gives information about the propagation time. The combination of matter traversed and a propagation time indicates the mean density through which the cosmic rays propagate. The general picture is that cosmic rays at these energies wander in a random walk fashion due to the random component in the galactic magnetic field, and that they spend a significant amount of time in the halo of the galaxy where the density of target particles is a great deal smaller than within the interstellar medium of the disk. A clear pedagogical treatment of the details can be found in Longair [12].

1.2.2 The knee and the ankle

The spectral index flattens relative to about 2.5 in the range $10^{14} - 10^{15}$ eV. Above this range the spectrum steepens fairly dramatically giving rise to the feature in the spectrum referred to as the knee. This steepening of the spectrum is often attributed to the diffusion of cosmic rays from the galaxy, their higher energy making them more difficult to contain.

In the range $10^{16} - 10^{19}$ eV, the spectral index steepens to about 3.08. A determination of the elemental composition in this range is difficult since detection is indirect,

fraught with large systematic errors, and blurred by model dependence. The data suggests a composition rich in heavy nuclei at the lower end of this range which becomes increasingly dominated by protons towards the high end. How this conclusion is reached is discussed in more detail in the last chapter of this work.

At energies above 10^{19} eV there is evidence for a flattening in the cosmic ray spectrum. The AGASA group finds a good fit with a broken energy spectrum characterized by a spectral index of about 3.16 ± 0.08 below $10^{19.01}$ eV and $2.78_{-0.33}^{+0.25}$ above [16]. This flattening of the spectrum is often referred to as the ankle.

1.2.3 The highest energy cosmic rays

There are only a handful of events observed with energies greater than 10^{20} eV. The most recent are six events recorded by the AGASA [4] ground array between 1990 and 1997, and an event recorded by Fly's Eye [3], an air fluorescence experiment. The highest energy event recorded by the AGASA array has an energy in the range of $(1.7-2.2) \times 10^{20}$ eV, while the Fly's Eye event was measured at $(3.2 \pm 0.9) \times 10^{20}$ eV (the highest energy recorded so far). Given the limited number of events, it is impossible to say anything with confidence concerning patterns in arrival directions or composition.

The GZK cutoff implies that the primaries associated with these events originated locally, within tens of megaparsecs (assuming they are protons). However there do not seem to be any accelerators in the local vicinity which seem capable of achieving these energies. It is tempting to try and extend Fermi's model of acceleration up to these energies, however particles can not attain arbitrarily high energy by this mechanism; as they attain higher and higher energies they become increasingly difficult to confine magnetically to the region of acceleration. The maximum energy obtainable in the simplest model of Fermi acceleration by a particle with charge Ze is given by

$$E_{\max} \approx \beta c \times Ze \times B \times L \quad (1.2)$$

where βc is the shock velocity, B is the magnetic field strength and L is the size of the acceleration region. It seems that only large structures associated with galaxies and groups of galaxies may potentially have the field strengths and dimensions needed

to achieve these energies. Such structures do not appear to lie within the requisite range.

It may be possible that some kind of compact object might accelerate particles to these energies with some kind of “one-shot” mechanism due to extremely intense electric fields. Closer examination suggests that coming up with a sufficiently efficient mechanism is nearly impossible since these types of environments also give rise to effects which degrade the energy of the accelerated particles.

Theorists are at the stage of entertaining any plausible scenario to explain the highest energy events. A later chapter in this work considers one such scenario: the highest energy cosmic rays are neutrinos which interact inelastically in the atmosphere in much the same way that nucleons do [17]. The problem of the GZK cutoff is avoided because the center of mass energies associated with the massless neutrino and cosmic microwave background photons are not high enough to generate pions. Such a suggestion requires unification of the fundamental forces at an energy much smaller than is conventional, however present limits from collider experiments do not rule out this possibility.

Is it possible that the events detected so far are not as energetic as they seem, or maybe they are something other than cosmic ray induced showers? Not likely. Firstly, AGASA and Fly’s Eye measure quantities closely related to the number of electrons in the shower front. A rough estimate of the number of electrons in the shower front during the bulk of its development is given by $\alpha \times E_0/\text{GeV}$ where α is a number of the order of (but less than) one. This means that the shower fronts for the highest energy events contain of the order of 10^{11} electrons. While fluctuations in this number are expected, it is difficult to overestimate energies by more than a factor of two. Secondly, outside of their enormous energies, the events measured by AGASA and Fly’s Eye are normal looking events. It is extremely unlikely that something other than a cosmic ray induced shower could coincidentally produce the characteristic signals. These issues are addressed in detail for the Fly’s Eye event [3].

Chapter 2

Electromagnetic Cascades

The longitudinal profile of a cosmic ray induced shower can be understood as a superposition of electromagnetic subshowers. These subshowers are initiated by the decay of neutral pions produced in the hadronic core of the shower. The general features of how this hadronic core evolves is detailed in the next chapter. For now it is enough to know that there is a mechanism by which the bulk of the energy in a proton or nucleus induced shower is ultimately transferred into neutral pions. Each of these neutral pions initiates a pure electromagnetic subshower by its decay into two photons. Consequently, a solid understanding of pure electromagnetic cascading is needed before progressing to showers containing a hadronic core. This chapter discusses general features of electromagnetic cascading and develops parameterizations describing pure electromagnetic cascades. Many of the features discussed here apply to hadronic showers as well; consequently, some attention is devoted to hadronic showers as a way of preparing for more detailed discussions to follow in later chapters.

The discussion centers around the standard formulas describing electromagnetic processes due to Bethe and Heitler [18]. This is entirely suitable for the energies at which parameterizations of pure electromagnetic cascades are required. Above about 100 PeV, the LPM effect, see Section 5.5, suppresses electromagnetic cross sections for interactions in air. Since the focus of this chapter is on the development of parameterizations at energies for which Bethe-Heitler physics is suitable, it is unnecessary to address the LPM effect at this stage.

2.1 Basic Description

There are two basic multiplicative processes which control the development of the longitudinal profile for pure electromagnetic cascades: pair production, the process by which a photon is destroyed and an electron-positron pair is created; and bremsstrahlung, where an electron or positron radiates a photon. At energies significantly greater than about 10^8 eV, these processes overwhelmingly dominate. These splitting processes can continue for many generations until the shower energy is dissipated through ionization losses by the electrons and positrons in the cascade.

The qualitative picture of the cascade is of a thin widening disc of particles propagating through the atmosphere at nearly the speed of light. As the disk propagates, the total number of particles grows exponentially as successive generations of particles lead to the generation of particles with lower and lower energies. The total energy of the particles in the disk at any stage plus the total of all energy losses up to that point which do not lead to the generation of further particles must add up to the energy of the primary particle which started the shower. While the disc widens due to scattering, the lateral spread can be considered separately from the longitudinal development of the shower. The discussion in this work focuses on a one-dimensional treatment of cascades.

Before getting into details, a brief comment on units is in order. All references to lengths and depths are understood to be in units of mass per area. This describes lengths in terms of the amount of matter traversed rather than in terms of distance. This is a natural way to describe cascading since electromagnetic processes unfold according to the amount of matter traversed. For media of a fixed density the two viewpoints of matter traversed versus distance are equal up to a proportionality constant, but in general the density of the medium can vary. For showers with small zenith angles, the depth at some stage in development, X , is related to its height in the atmosphere, h , through the relation,

$$X = \sec(\theta) \times \int_h^\infty \rho(h')dh'. \quad (2.1)$$

Here, $\rho(h)$ is the density profile of the atmosphere and θ is the zenith angle of the

shower (showers coming in at larger zenith angles traverse more matter to reach a fixed height than showers coming in closer to the vertical). Decay processes such as $\pi^+ \rightarrow \mu + \nu_\mu$ are best described in units of distance (taking all particles to be moving at approximately the speed of light). In such cases, Monte Carlo simulation requires changing between these two viewpoints so that a random sampling between decay and interaction is possible.

At an elementary level, pair production and bremsstrahlung can be viewed as splitting processes in which the initial particle splits into two after propagating some length. Pair production and bremsstrahlung are closely related processes, and consequently they are described by the same length scale. All such splitting processes give rise to the same general features when applied to cascading. Thus for the sake of simplicity, consider a model where the distance of propagation in g/cm^2 is a fixed length for all particles and that the energy of a parent particle is equally divided between its two offsprings. This toy model was first discussed by Heitler in an elementary treatment of electromagnetic cascading [19]. The number of particles in the cascade as a function of depth is given by

$$N(X) = 2^{X/\lambda}, \quad (2.2)$$

where X is the depth in the atmosphere and λ is the propagation length for the splitting process. The cascade continues to grow until the particles reach a critical energy, E_{crit} , where ionization losses become the dominant process at which point the particles are quickly absorbed. At this depth, the shower has reached its maximum size which is

$$N(X_{\text{max}}) = E_0/E_{\text{crit}} \quad (2.3)$$

where E_0 is energy of the initiating particle. Solving for the depth at which shower maximum is reached,

$$X_{\text{max}} = \lambda \frac{\ln(E_0/E_{\text{crit}})}{\ln 2}. \quad (2.4)$$

Obviously this is a highly simplified picture: propagation and splitting are random processes and so qualitatively it is expected that the shower gradually diminishes in

size once it passes maximum. Still, some basic conclusions can be drawn:

$$N_{\max} \propto E_0 \quad (2.5)$$

and

$$X_{\max} \propto \ln(E_0/E_{\text{crit}}). \quad (2.6)$$

Both are general properties of multiplicative cascades and also hold true for hadronic showers. These features carry over to actual electromagnetic cascades because the products of pair production and bremsstrahlung have energies less than but of the order of the parent energy while losses due to ionization are nearly constant exhibiting only a mild logarithmic dependence on the electron energy (assuming the electrons are relativistic, $E \gg mc^2$). Thus, ionization losses are inconsequential until losses over the characteristic length are of the order of the particle's energy.

2.2 The Track Length Integral

An important feature of cascades is the track length integral which relates the longitudinal profile to the energy of the initiating particle,

$$E_0 \approx \alpha \times \int N_e(X) dX. \quad (2.7)$$

Here, $N_e(X)$ is the number of electrons and positrons in the shower at depth X and α is the average energy loss rate per particle. Technically, when this is extended to hadronic cascades the integral is over all charged particles; however, the overwhelming majority are electrons and positrons. Also for hadron initiated showers, a fraction of the shower's energy (about 5% for protons around 10^{20} eV and 10% for iron at the same energy) is channelled into muons and neutrinos. In these cases the track length integral gives a lower bound on the primary energy. This is demonstrated in Chapter 4 where this fractional energy loss is directly calculated.

Since the ionization energy loss rate is characterized by a mild logarithmic dependence on the particle energy, the average energy loss rate can be taken at the energy at which the majority of electrons and positrons are found. Keeping in mind

the simple splitting model of electromagnetic cascades, this means energies at which ionization losses represent the dominant process ($E \leq E_{\text{crit}}$). There are a number of formal definitions of the critical energy. Often it is defined as the energy at which energy losses due to bremsstrahlung equal those due to ionization, or sometimes as the energy at which the ionization losses per radiation length are equal to the electron's energy [20]. The radiation length, denoted here as X_{rad} , is the natural length which characterizes electromagnetic cascading and is defined as the length over which a high energy electron loses all but $1/e$ of its energy [21] ($X_{\text{rad}} \approx 37.1\text{g/cm}^2$ for electrons in air). Approximately, both definitions of the critical energy give the same numerical value, $E_{\text{crit}} = 81\text{MeV}$. The energy loss rate due to ionization is then given by

$$\alpha \approx \frac{E_{\text{crit}}}{X_0} = 2.18\text{MeV}/(\text{g/cm}^2). \quad (2.8)$$

Detailed simulations by Hillas [22] indicate that properly including the energy dependence of ionization rates and integrating over the energy distribution of particles leads to an average energy loss per particle of $2.24\text{MeV}/(\text{g/cm}^2)$, a small difference given the large errors present in the reconstruction of the longitudinal profile from an experimental signal. Since this energy loss rate is characteristic of the region near the critical energy, it will not significantly change as a function of the primary energy for ultra-high energy cascades since the primary energy is many orders of magnitude higher than the critical energy. Consequently, the choice of energy loss rate effects only the overall normalization of the longitudinal profile, not its shape.

Another consequence of the primary energy being many orders of magnitude greater than the critical energy is that the correspondence between the track length integral and the energy deposited in the electromagnetic cascade is not subject to significant stochastic fluctuations over sets of showers characterized by the same primary energy. (for a shower of energy 10^{20}eV , the number of particles at shower maximum is of the order of 100 billion). Large fluctuations in the shape of the profile can be traced back to the early stages of the cascade where only a few particles carry most of the shower's energy, however these fluctuations do not effect the area underneath the profile. Area is fixed by the later stages of shower development. With respect to showers induced by protons or nuclei, the stochastic nature of charged pion decay

versus interaction can introduce large fluctuations in the fractional energy channelled into the electromagnetic cascades for primary energies around a PeV. But for ultra-high energies greater than 10^{18} eV virtually all charged pions interact. In fact, at these energies it is important to consider the hadronic interaction of neutral pions.

Given the large number of particles which make up a shower, Simulation must make use of parameterizations which describe the average development of subshowers. The general technique is to express the longitudinal profile as a superposition of many subshowers each of which contributes only a small part to the track length integral. Each of these subshowers can then be parameterized by its average without compromising the modeling of fluctuations in the development of the total shower. This is done by parameterizing all particles with energies below an energy threshold, $E_{\text{thr}} = x_{\text{thr}} \times E_0$ with $x_{\text{thr}} \ll 1$. All parameterizations which are developed for this purpose are constrained such that the track length integral gives the energy channelled into the electromagnetic cascade.

2.3 Analytic Cascade Theory

What are sought are solutions to cascade equations which represent the average development of longitudinal profiles corresponding to a single initiating particle. The classic works which remain the standard reference on analytic cascade theory are due to Rossi and Greisen [21, 24]. A useful summary appears in Gaisser's book [14]. The equations to be solved are coupled cascade equations which describe the development of photons and electrons (together with positrons). Analytic calculations of this type are prohibitively difficult unless some approximations are made. Even with approximations, exact solutions are unattainable.

There are two standard approximations that are typically made. What is referred to as approximation A treats pair production and bremsstrahlung in the "complete screening" limit. In this high energy limit the atomic electrons screen the nucleus so that these processes occur only for impact parameters less than the effective radius of the atom. For electrons in air, this holds for energies greater than about 40 MeV. In this approximation, all cross sections scale with energy. In other words, the prob-

ability that a parent particle of energy E_0 after traversing a distance dX produces a particle of energy between vE_0 and $(v + dv)E_0$ is a function of only v and contains no reference to the parent energy. Ionization losses are ignored. Approximation B is the same except that ionization losses are included as a constant energy loss term.

To motivate the form of the parameterizations which are adopted, consider a cascade equation,

$$\frac{dn(E, t)}{dt} = -n(E, t) + 2 \int_0^1 n\left(\frac{E}{v}, t\right) \frac{dv}{v}, \quad (2.9)$$

which gives the differential energy spectrum of particles, n , as a function of depth, t . It is explicitly in scaling form. This refinement of the simple splitting model discussed above describes propagation with the depth expressed in units of the mean free path and splitting into two particles of the same type with a flat distribution between zero and the parent energy. In terms of a Monte Carlo algorithm, this corresponds to choosing the distance of propagation by sampling an exponential distribution with unit mean free path and splitting the parent energy using a uniform deviate in the range zero to one to sample the fractional energy of one of the two secondary particles with the remaining fraction going to the other. Though this cascade equation is simple, its solution entails the same procedures as in calculations under Approximation A.

To solve this equation for the initial condition of a single particle, begin by expressing the cascade equation in terms of the Mellin transform of $n(E, t)$,

$$M(s, t) = \int_0^\infty E^s n(E, t) dE. \quad (2.10)$$

This is done by multiplying the cascade equation by E^s and integrating over E . After a little manipulation the transformed equation is,

$$\frac{dM(s, t)}{dt} = \lambda(s)M(s, t). \quad (2.11)$$

where $\lambda(s) = (-1 + 2/(s+1))$. Applying the boundary condition $n(E, 0) = \delta(E - E_0)$, the solution to this equation is

$$M(s, t) = E_0^s \exp[\lambda(s)t]. \quad (2.12)$$

As with the analogous equations for photons and electrons under approximation A, it is only in the inversion of the Mellin transform of $n(E, t)$ that approximations are required,

$$n(E, t) = \frac{1}{2\pi i} \int_{-i\infty+s_0}^{+i\infty+s_0} E^{-(s+1)} M(s, t) ds. \quad (2.13)$$

After substituting in $M(s, t)$ and with a little rearrangement, the solution takes the form,

$$n(E, t) = \frac{1}{E} \frac{1}{2\pi i} \int_{-i\infty+s_0}^{+i\infty+s_0} \exp[\lambda(s)t + sy] ds, \quad (2.14)$$

where $y = \ln(E_0/E)$. The integral can be approximated using the method of steepest descents. This involves expanding the argument of the exponent up to second order in s about its saddle point leaving a Gaussian integral to be computed by contour integration. The result is,

$$n(E, t) \approx \frac{1}{E \sqrt{2\pi \lambda''(s)}} \exp[\lambda(s)t + sy] \quad (2.15)$$

where s is a function of the depth t and is determined by the relation

$$\lambda'(s) + y = 0 \quad (2.16)$$

which is the condition for the extremum in the argument of the inverse Mellin transform integral.

Approximately, the depth at maximum, T , is reached when the argument of the exponential attains a maximum,

$$\left\{ \frac{d}{ds} [\lambda(s)t + sy] \right\} \frac{ds}{dt} + \lambda(s) = \lambda(s) = 0. \quad (2.17)$$

In general, for cascade equations of the type under consideration this is satisfied when $s = 1$ as can be verified for the present case. The parameter s is called the *age parameter* for the reason that the shower grows in size as s grows until s reaches one after which the shower decreases in size. Solving for T ,

$$T = 2y = 2 \ln \frac{E_0}{E}. \quad (2.18)$$

The depth at maximum has the expected functional dependence on the primary and target energies.

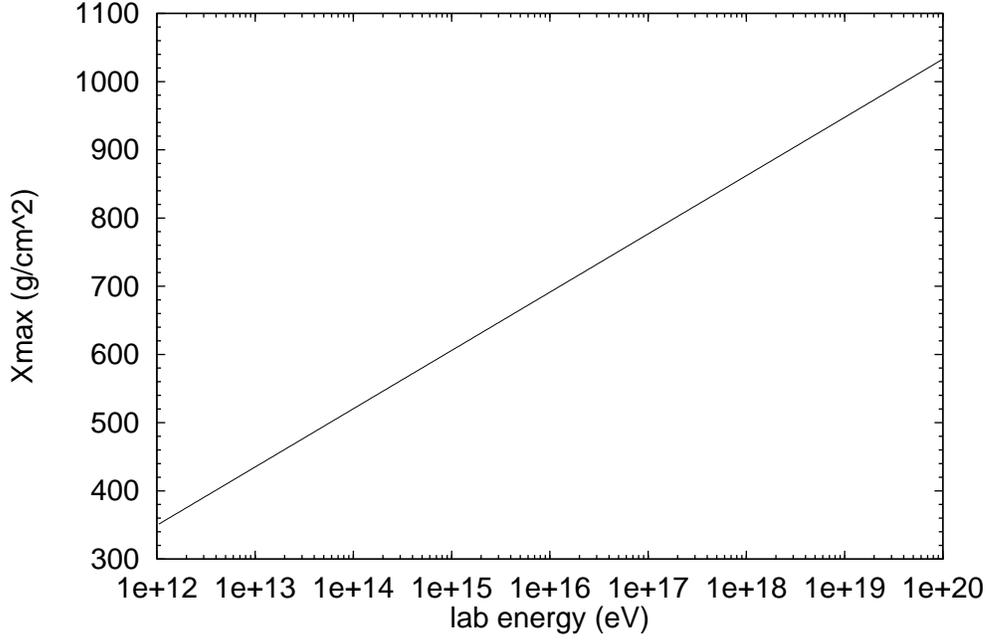


Figure 2.1: The depth at maximum as a function of primary photon energy as prescribed by standard electromagnetic cascade theory.

2.4 The Greisen Formula

An expression introduced by Greisen [25] which describes the total number of electrons and positrons in a photon induced cascade under approximation B is given by,

$$N_e(t) \approx \frac{0.31}{\sqrt{y}} \exp[\lambda_1(s)t + sy] \quad (2.19)$$

where t is the depth measured in radiation lengths,

$$y = E_0/E_{\text{crit}}, \quad (2.20)$$

$$\lambda_1 \approx \frac{1}{2}(s - 1 - 3 \ln s) \quad (2.21)$$

and

$$s = \frac{3t}{t + 2y}. \quad (2.22)$$

The function $\lambda_1(s)$ emerges in the same manner as $\lambda(s)$ in the simplified cascade equation. Note the similarity to the solution for the simplified case. Solving for

$s = 1$, the depth at maximum is given by

$$T \approx \ln \frac{E_0}{E_{\text{crit}}}. \quad (2.23)$$

Figure 2.1 shows the depth at maximum as a function of photon energy where the depth has been converted to units of g/cm^2 . The slope of this line is called the elongation rate. That the elongation rate is constant is a reflection of the scaling nature of the physics described so far. Any scaling model of physics which ultimately reduces the cascade to a superposition of photon induced subshowers is characterized by this same elongation rate. Experimental measurements of a changing elongation rate can be a tell-tale sign of a changing composition or of scaling violations in the physics [13]¹.

Since the solution under approximation B includes ionization losses as a constant energy loss term, the overall constant in front of the exponent in Equation 2.19 can be understood as a normalization factor which gives the correct value for the track length integral. With some manipulation, the solution can be expressed in a more conventional form,

$$N_e(t) = \frac{0.31}{\sqrt{y}} \exp\left[t\left(1 - \frac{3}{2} \ln s\right)\right]. \quad (2.24)$$

Recent Monte Carlo simulations which treat more detail than is possible with analytic methods verify that this formula describes photon induced showers quite well [26].

Figure 2.2 shows the Greisen formula for a photon shower of energy 10^{20}eV . The basic shape of the profile is characteristic of all longitudinal profiles including those induced by hadrons: the area is fixed by the energy, and the depth at maximum lies near $\ln(E/E_{\text{crit}})$. An important feature to notice is how deep the bulk of the shower develops. The units of depth have been translated to units of g/cm^2 . The length which characterizes the early interactions in the shower is the radiation length, $X_{\text{rad}} = 37.1\text{g}/\text{cm}^2$. It is the early interactions which are the only source of significant fluctuations in shower development. The length associated with fluctuations is much

¹Recall that scaling refers to the fact that no explicit reference to the primary energy is needed to describe bremsstrahlung or pair production. With respect to hadronic interactions, when there are effects which require knowledge of the energy to describe interactions, the physics is said to exhibit scaling violations.

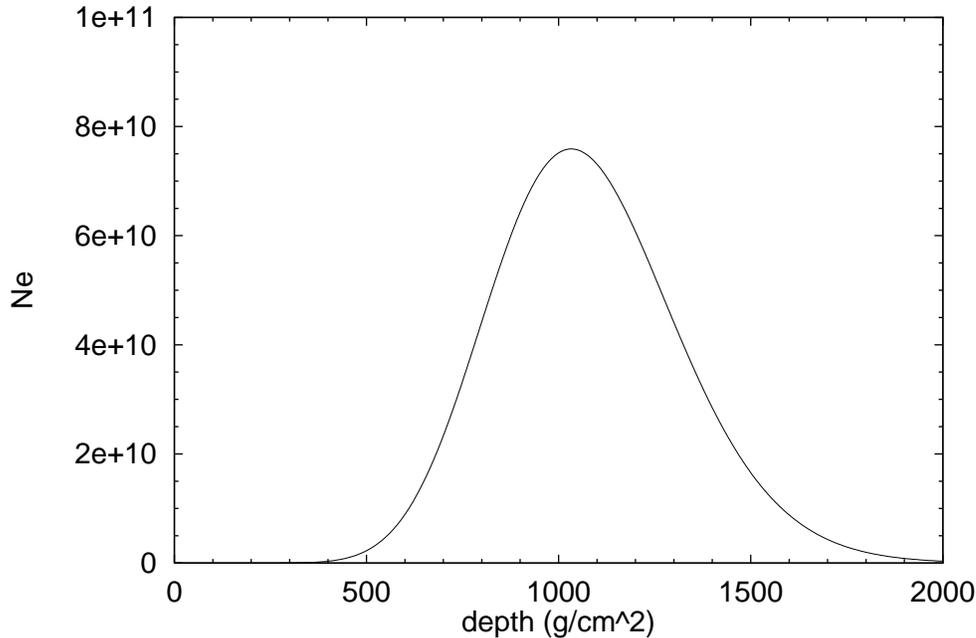


Figure 2.2: Shower size versus atmospheric depth as described by the standard Greisen formula for a 10^{20} eV photon initiated shower.

less than the depth at maximum of the total shower which is around $1000\text{g}/\text{cm}^2$. Therefore fluctuations in the development of the longitudinal profile consist of small deviations from the average in the depth at maximum, X_{max} , and in the size at maximum, N_{max} .

The same holds true for hadronic showers where the length scale of interactions is of the order of the radiation length, though the rise to maximum is somewhat steeper giving a shallower depth at maximum by about $100\text{g}/\text{cm}^2$ for proton induced showers and about $200\text{g}/\text{cm}^2$ for iron induced showers. The steeper rise is due to the large number of pions that can be produced in a single hadronic interaction. The number is of the order of hundreds for hadrons on a nuclear target. For nuclear primaries there is the added effect that the shower begins in a more advanced stage of development since it is comprised of a number of nucleons which share the primary energy equally. Still the basic arguments concerning electromagnetic cascades apply here as well. It is safe to say that unless there are dramatic changes in the lengths associated with the early stages of shower development, the characteristic shape depicted in

Figure 2.2 holds true for all showers at the highest energies even once fluctuations in development are considered.

2.5 A Modified Greisen Formula

Equation 2.24 describes photon induced cascades relative to the photon's depth of creation. However all pure electromagnetic showers, such as those induced by electrons or the decay of neutral pions, evolve in a similar way. While the depth at maximum for showers of a fixed energy may differ slightly depending on the nature of the initial conditions, all pure electromagnetic showers scale in the same manner as the primary energy is varied. In other words the depth at maximum in radiation lengths for any pure electromagnetic shower is given by,

$$T \approx \ln \frac{E_0}{E_{\text{crit}}} + \beta, \quad (2.25)$$

where the constant β depends on the initial conditions of the shower. The scaling with energy is a reflection of the multiplicative processes which govern the development of the shower (pair production and bremsstrahlung) irrespective of the shower's initial conditions. This suggests that the average profile for any type of electromagnetic shower can be fitted with a modified Greisen formula of the form,

$$N_e(t - \Delta) = A \exp\left[t\left(1 - \frac{3}{2} \ln(s)\right)\right], \quad (2.26)$$

where

$$s = \frac{3t}{t + 2y} \quad (2.27)$$

with

$$y = E_0/E_{\text{crit}} - \alpha. \quad (2.28)$$

A positive Δ pushes the entire profile to a deeper depth without changing the shape while a positive α gives a shallower depth at maximum by steepening the rise in the profile. The factor A is fixed by requiring the track length integral to give the total subshower energy. The factors Δ and α can be determined for a variety of initial conditions with the help of Monte Carlo simulation of the earliest stages of shower

development and requiring consistency with the prescription $\Delta = 0$ and $\alpha = 0$ for photon induced cascades relative to the photon's depth of creation.

A form which only considers a shift of the entire profile was used to fit a detailed simulation of electron induced showers [26]. The additional parameter α included here can be meaningfully fitted because the approach adopted here is different: the focus is to develop parameterizations which imply the highest level of consistency with the standard Greisen formula which describes photons. The techniques described below sacrifice a precise modeling of bremsstrahlung in exchange for a higher level of consistency between Monte Carlo simulation and parameterization. Consequently, there is negligible error associated with average quantities such as the average depth at maximum, but potentially small errors in the description of fluctuations. These would be evident only in pure electromagnetic cascades where early development contains electrons which individually carry a significant proportion of the total shower energy. In the Fly's Eye analysis of the mass composition [27] (which predominantly relies on the elongation rate, an average quantity), systematic errors associated with the hybrid use of the Greisen formula along with Monte Carlo techniques are estimated at about $10\text{g}/\text{cm}^2$. While small with respect to the reconstruction errors associated with individual showers, it is important in the analysis of sets of showers.

The prescription $\alpha = 0$ and $\Delta = 0$ is essentially the standard Greisen formula, though the overall normalization is adjusted by a few percent to give the appropriate contribution to the track length integral. Ultimately this difference is negligible compared to other errors associated with an experimental measurement. However, in a practical sense it is useful to enforce the relationship between energy and the track length integral precisely since the energy channelled into the electromagnetic cascade is a quantity which is easily tracked in a Monte Carlo simulation. In any case, an exact calculation under approximation B, if it were possible, would respect this relation exactly since energy loss is only through a constant term. Tracking the energy deposited into the electromagnetic cascade separately from the actual profile aids in the fitting of profiles: the fitted profile can be constrained to give the appropriate track length integral.

2.6 Procedure for developing parameterizations

The remainder of this chapter is concerned with fixing parameters for a number of initial conditions. An average profile is constructed through a combination of Monte Carlo simulation and parameterization. Beginning with the Monte Carlo simulation of propagation and/or generation, particles are parameterized as soon as they are in a form for which the parameterization is known. Average profiles are generated for a set of 10,000 primaries of fixed type and energy, all injected at a depth of zero. A fit is performed by adjusting the modified Gresien parameters, α and Δ , to maximize the area of overlap between the fitted and the actual profiles. This is chosen over a least-squares fitting procedure since it is more closely connected with the track length integral which is measured experimentally. This criterion naturally suppresses fitting the shower in regions where the shower size is small.

With 10,000 trials, the errors are reduced to about 0.01 radiation lengths in $\Delta - \alpha$ (the total shift in the depth at maximum) and 0.1 radiation lengths in the correlated values of the Δ and α . In other words, the fit is sensitive to the overall shift in the depth of maximum but rather insensitive to whether the shift is accomplished through a slight deformation in shape or through a shift of the entire profile. The discrepancy between the simulated and the best fit profiles is generally of the order 0.1% or less.

While the modified Greisen formula scales with energy, it is still best to conduct the simulations at an appropriate energy. The goal is to understand the features of showers of energies greater than 10^{20} eV. So that fluctuations are captured, parameterizations are used only for particles with energies less than a specified simulation threshold. An energy threshold of $1/1000 \times E_0$ is suitable. This can be seen by comparing fluctuations in proton initiated showers as a function of threshold. The result is that any threshold lower than about $1/300 \times E_0$ sufficiently captures fluctuations. Intuitively, it is easy to see that this is the case since each subshower contributes to the track length integral in proportion to its energy. Thus an appropriate energy for the purpose of developing parameterizations is about 10^{17} eV. This is used as the initial energy for all simulations discussed in the following sections.

2.7 Neutral pions

For showers induced by protons or nuclei, the electromagnetic component of the shower is fed by the decay of neutral pions produced in the hadronic core of the shower. For applications in which the average properties of full hadronic showers are sought, it is sufficient to understand electromagnetic cascading only at the level of the average profiles of subshowers induced by the decay of neutral pions.

In the rest frame of the neutral pion, the photons share the energy of the pion equally with their momenta directed opposite one another. When boosted to the lab frame, the energy of the photon oriented at angle θ with respect to the boost axis is given by

$$E = \frac{m_\pi}{2}\gamma(1 + \beta \cos \theta) \quad (2.29)$$

where γ is the Lorentz factor and β is the velocity of the neutral pion in the lab frame. Since the decay is isotropic in the lab frame, the distribution of photons is

$$\frac{dn}{d\Omega} \propto \frac{dn}{d \cos(\theta)} \propto \frac{dn}{dE} = \text{constant}. \quad (2.30)$$

In the ultra-relativistic limit where $\beta \approx 1$, this implies that the energy distribution of photons in the lab frame is flat between zero and the parent neutral pion's energy. The simulation is performed by assigning the fractional energy of one of the photons as a uniform deviate and assigning the remaining fraction to the other photon. These photons are parameterized using the modified Greisen formula with $\alpha = 0$ and $\Delta = 0$.

The resulting fit for the modified Greisen formula which describes neutral pions relative to their depth of decay is $\alpha = 0.33$ and $\Delta = -0.17$. The overall shift in the depth at maximum, $\Delta - \alpha = -0.50$, can be understood with a simple calculation. The photons from the decay of the neutral pion contribute to the profile each with a weight proportional to its energy. Thus an approximation to the final depth at maximum is a weighted average over the depths at maximum in the distribution of photons stemming from the decay. Using Equations 2.23 and 2.30,

$$T \approx \int_0^1 \{x \ln[xE/E_{\text{crit}}] + (1-x) \ln[(1-x)E/E_{\text{crit}}]\} dx. \quad (2.31)$$

Here x represents the fractional energy of one of the photons relative to the parent neutral pion. All values of x are equally probable. Carrying out the integration,

$$T \approx \ln(E/E_{\text{crit}}) - 0.5 \quad (2.32)$$

which indicates that the average profile for neutral pions reaches maximum 0.5 radiation lengths sooner than a photon shower of the same total energy.

2.8 Deconvoluted photons

The standard Greisen formula describes photon induced subshowers relative to the photon's depth of creation. It is desirable to have a parameterization which describes the subshower relative to the photon's depth of pair creation since the sampling of starting point fluctuations is much less computationally intensive than the parameterization of the subshower. Starting point fluctuations are the dominant source of fluctuations in development. A first approximation pulls the standard Greisen formula back by one mean free path and pushes it forward by randomly sampling an exponential distribution. In this approximation, the shape of the profile still reflects the convolution over starting depths. A better method is to find those values of α and Δ which when convoluted over starting point fluctuations implies the average profile relative to the photon's depth of creation. The shape of the resulting profile is slightly steeper (giving a slightly larger size at maximum) to counteract the mild flattening which occurs from convoluting over starting depths.

The mean free path for pair creation under approximation A in units of radiation lengths is given by

$$\lambda_{\text{pair}} = \left(\int_0^1 \psi(u) du \right)^{-1} \approx 1.29, \quad (2.33)$$

where $\psi(u)dtdu$ is the probability that after propagating dt radiation lengths, a photon produces a pair in which the positron (or equivalently the electron) carries a fraction u of the photon's energy. The explicit form of the function $\psi(u)$ is needed in the next section when electron induced subshowers are addressed, but for now all that is needed is the numerical value for the mean free path.

For photons relative to the depth of pair creation, $\alpha = 1.10$ and $\Delta = -0.19$ gives the best fit. As expected, $\alpha - \Delta = 1.29$ gives the mean free path, and the bulk of the shift is accomplished through α which implies a slightly steeper profile.

2.9 Electrons

The results of the last section can be used in the description of subshowers induced by electrons relative to their creation depth. The simulation proceeds by carrying out pair production and parameterizing the resulting leptons. The profile constructed after summing many events should give the average profile for photons relative to their depth of pair creation.

To carry out pair production, the energy of one of the leptons is sampled using the probability distribution [14],

$$\psi(u) = \frac{2}{3} - \frac{1}{2}b + \left(\frac{4}{3} + 2b\right) \left(u - \frac{1}{2}\right)^2 \quad (2.34)$$

where u is the fractional energy of the lepton and,

$$b = (18 \ln[183/Z^{1/2}])^{-1} \approx 0.0122. \quad (2.35)$$

The distribution is symmetric about $u = 1/2$, so there is no difference between the electron and the positron as far as cascading is concerned. The search through parameter space gives $\alpha = 0.80$ and $\Delta = -0.01$ for electrons relative to their depth of creation.

2.10 Bremsstrahlung and an effective model

In approximation A, the probability for an electron (or positron) after propagating dt radiation lengths to radiate a photon which carries off a fraction v of the electron's energy is given by [14],

$$\phi(v) = v + \frac{1-v}{v} \left(\frac{4}{3} + 2b\right). \quad (2.36)$$

In the context of a Monte Carlo algorithm, this expression is difficult to handle because it is logarithmically divergent as $v \rightarrow 0$. In traversing a depth dt , an infinite number of photons with energies arbitrarily close to zero are radiated,

$$\frac{dN(v < \delta)}{dt} = \int_0^\delta \phi(v)dv = \infty. \quad (2.37)$$

Consequently, it is impossible to assign a mean free path for the radiation of a single photon even though the energy loss rate is perfectly well defined,

$$\frac{dE}{dt} = \int_0^1 (vE)\phi(v)dv = -E(1 + b) \approx -E. \quad (2.38)$$

The standard procedure for dealing with this problem is to introduce a cutoff, v_{\min} , which is much less than the smallest energy of interest. Above this cutoff, a Monte Carlo splitting algorithm can be used though care must still be taken to treat the energy losses below this cutoff [14]. Here, a much simpler approach is adopted. For showers induced by protons or nuclei, it is rare that an electron is produced that carries a significant fraction of the shower's energy. Consequently, any Monte Carlo algorithm that introduces fluctuations in electron subshowers in some manner while still implying the correct average profile is suitable for most purposes. This approach which focuses on self consistency with the parameterization for photon induced cascades may even be preferable since it avoids the systematic error in the depth at maximum that can result from applying a hybrid Monte Carlo/parameterization scheme to the cascading of photons and electrons as discussed above.

2.10.1 A simple effective splitting model

A simple model which overestimates fluctuations is an effective splitting model with a uniform deviate and with an associated effective mean free path. Following the same line of argument that led to Equation 2.31, the effective mean free path satisfies,

$$\ln(E/E_{\text{crit}}) = \lambda_{\text{eff}} + \int_0^1 \{x \ln[xE/E_{\text{crit}}] + (1-x) \ln[(1-x)E/E_{\text{crit}}]\} dx. \quad (2.39)$$

With a little manipulation this reduces to,

$$\lambda_{\text{eff}} = -2 \int_0^1 x \ln(x) dx = 0.5. \quad (2.40)$$

This model ignores the distinction between leptons and photons and considers an effective particle which splits to create particles of the same type.

2.10.2 A more general effective splitting model

This model can be made more general by introducing a parameter which controls the splitting. Suppose the fractional energy of one of the resulting particles is drawn from the distribution

$$p(x)dx = (\beta + 1)(1 - x)^\beta, \quad (2.41)$$

with the remaining fractional energy going to the other particle. The generalization of Equation 2.40 takes the form,

$$\lambda_{\text{eff}} = - \int_0^1 (\beta + 1)(1 - x)^\beta [x \ln(x) + (1 - x) \ln(1 - x)] dx. \quad (2.42)$$

The average fractional energy of the particle sampled from Equation 2.41 is given by

$$\langle x \rangle = \int_0^1 x(\beta + 1)(1 - x)^\beta dx = \frac{1}{2 + \beta}. \quad (2.43)$$

So the general splitting model requires the setting of a single parameter which can be expressed either as an effective mean free path or as the average energy of the typically lower energy effective particle. Table 2.1 gives the results of Monte Carlo simulations for a number of effective mean free paths. As expected, the difference between $\alpha - \Delta$ for electrons relative to the depth of creation and for the effective particles relative to the depth of splitting is equal to the effective mean free path.

After deconvoluting for an effective mean free path, the splitting algorithm can be used to fix the average energy of the typically lower energy particle, $\langle x \rangle$. The particles resulting from the split are immediately parameterized using the average profile for electrons relative to their depth of creation. The profile constructed in this way is compared with the modified Greisen formula describing the deconvoluted effective particle. In other words, once an electron is created, it is viewed as an effective electromagnetic cascading particle. Since α and Δ are already fixed, $\langle x \rangle$ is the only free parameter. This parameter has been fixed in two independent ways: by the Monte Carlo technique as used above and through Equation 2.42. Both methods

λ_{eff}	α	Δ	error (%)	$\langle x \rangle$	error (%)
0.5	0.97	-0.34	0.06	0.50	0.39
0.4	0.91	-0.30	0.04	0.18	0.17
0.3	0.86	-0.25	0.02	0.11	0.06
0.2	0.83	-0.18	0.02	0.024	0.16
0.1	0.81	-0.10	0.02	0.060	0.21
0.0	0.80	-0.01	0	0	0

Table 2.1: Modified Greisen parameters and the average fractional energy of the typically lower energy particle for effective electromagnetic splitting models characterized by an effective mean free path. The error represents the discrepancy in area overlap between the simulated and the target profiles.

agree up to two significant figures which is all that is allowed by the Monte Carlo technique with 10,000 iterations; in any event, Equation 2.42 is only approximate. Table 2.1 shows the results.

2.10.3 Choosing a model

Which effective model is most appropriate? The simplest answer is to choose a middle of the road model which does not greatly overestimate fluctuations but doesn't ignore them either. The effective splitting model with $\lambda_{\text{eff}} = 0.5$ certainly overestimates fluctuations since it samples a flat distribution for the splitting. Adopting $\lambda_{\text{eff}} = 0.3$ is probably a reasonable choice; it also has the benefit of having the lowest error of the models investigated so far (see Table 2.1).

The errors listed in Table 2.1 suggest a method for choosing an appropriate model. One might expect that the error would decrease as the limit of $\lambda_{\text{eff}} = 0$ and $\langle x \rangle = 0$ is approached. However this is not the case, the errors for the models investigated suggest a minimum exists. Though all the errors are fairly small, they reflect the degree to which an effective model if carried out in full gives the profile shape which is characteristic of electromagnetic cascading. While any effective model does a good job of representing the average depth at maximum as a function of energy, the shape of the profile is sensitive to the choice of model. To illustrate this, consider the effective model with $\lambda_{\text{eff}} = 0.5$. Figure 2.3 compares the modified Greisen formula

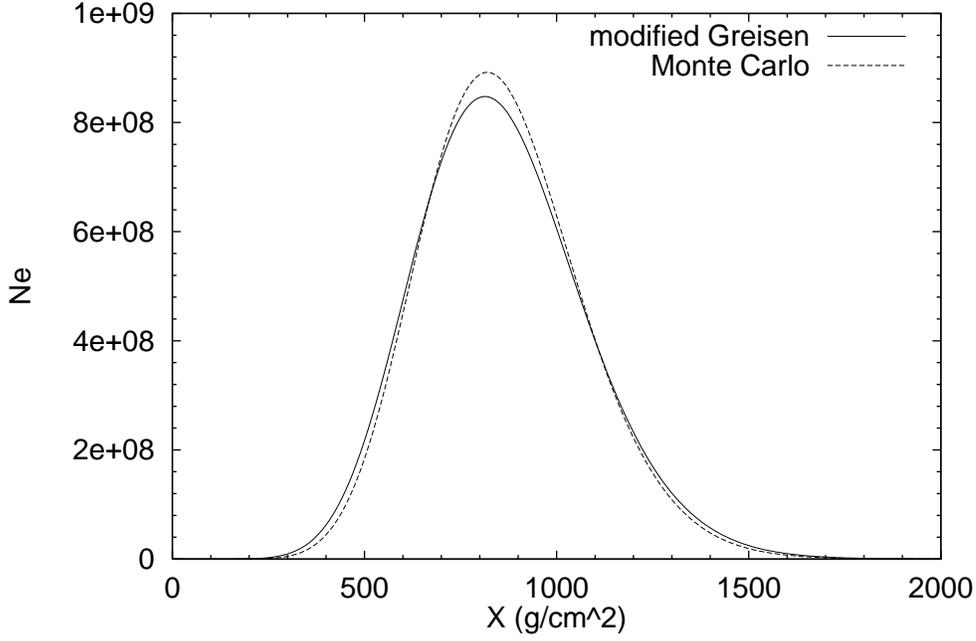


Figure 2.3: Comparison of the longitudinal profiles for primaries at 10^{18} eV using an effective splitting model with $\lambda_{\text{eff}} = 0.5$. The Monte Carlo averaged 1000 showers using a simulation threshold of 10^{15} eV.

for electrons relative to their depth of splitting to an average profile constructed by simulating with the effective model. The primary energy is 10^{18} eV. For the Monte Carlo, particles below 10^{15} eV were parameterized and the average was constructed over 1000 showers. While the depths at maximum agree, the sizes at maximum differ by close to ten percent.

A random search was performed to find the most appropriate splitting model. The correspondence between λ_{eff} and α as shown in Table 2.1 is well represented by a power law,

$$\alpha = 0.657(\lambda_{\text{eff}})^{1.95}. \quad (2.44)$$

Once α is fixed, Δ is determined through the relation,

$$\lambda_{\text{eff}} = \alpha - \Delta \quad (2.45)$$

so that the correct depth at maximum results; $\langle x \rangle$ is fixed by Equation 2.42 where it has been tabulated so that interpolation can be used to solve for $\langle x \rangle$

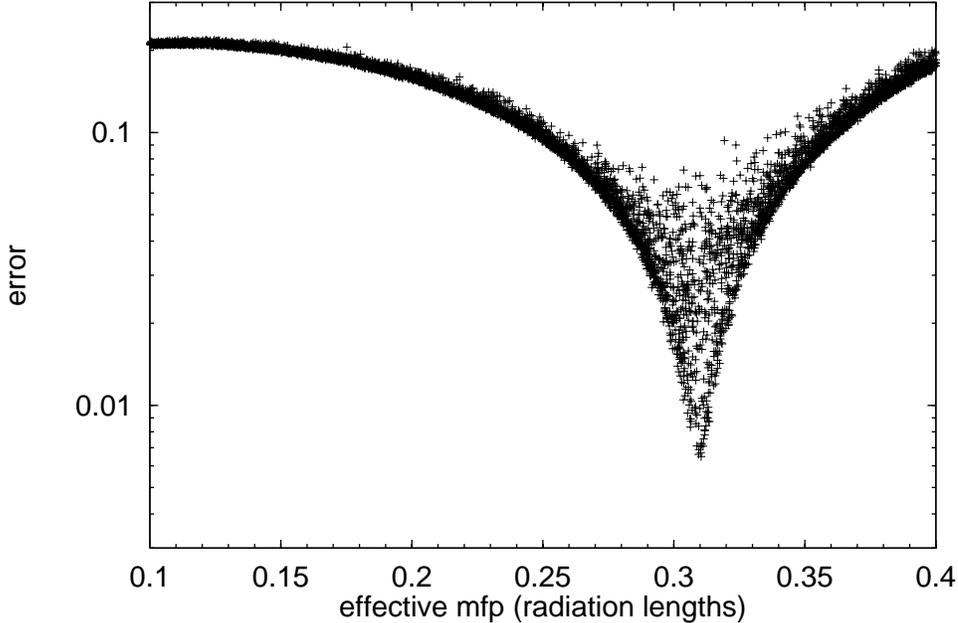


Figure 2.4: Results of the random search for the best effective splitting model. There is a clear minimum near $\lambda_{\text{eff}} = 0.31$.

numerically. The result is shown in Figure 2.4; a sharp minimum is clearly evident. The profiles depicted in Figure 2.3 were recalculated and are shown in Figure 2.5. The near perfect agreement suggests that this effective model adequately describes electromagnetic cascading. This result also confirms that the modified Greisen parameters can be taken to be energy independent since the simulation which utilized the effective splitting model parameterized electrons at energies orders of magnitude lower than the energy at which the modified Greisen parameters were calculated.

2.11 Fluctuations in shower development

To what degree does the use of an effective model for electrons compromise the modeling of fluctuations in hadronic showers? The dominant source of fluctuations in hadronic showers is in the starting depth of the shower (the depth of the first hadronic interaction), and to a lesser extent development fluctuations associated with the few highest energy secondaries with energies comparable to that of the prima-

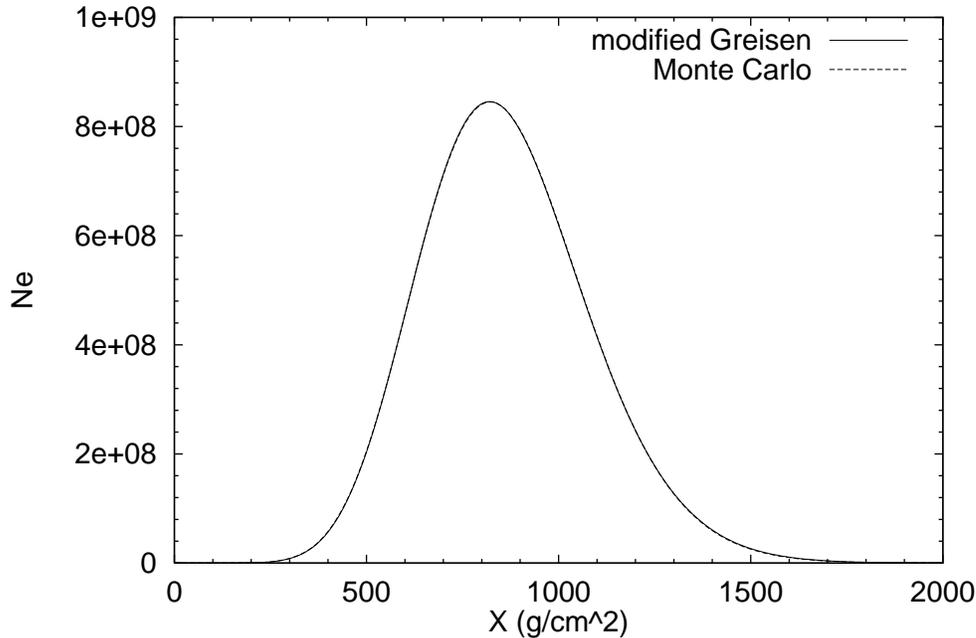


Figure 2.5: Comparison of the longitudinal profiles for primaries at 10^{18} eV using the best effective splitting model. The Monte Carlo averaged 1000 showers using a simulation threshold of 10^{15} eV.

ry. Development fluctuations associated with electrons is of little consequence since electrons do not typically possess energies comparable to the primary energy at any stage in shower development. This is due to the multistage process which leads to their generation.

The dominance of starting point fluctuations is illustrated in Figure 2.6 which shows the depth at maximum distribution for photon induced cascades at 10^{18} eV. The effective splitting model is compared with immediate parameterization of electrons. In the case where electrons are parameterized immediately, only the primary photon's depth of pair production along with the energy assignment of the resulting electrons are sampled randomly. When an effective splitting model is employed the width of the distribution widens slightly and a shallow depth tail emerges for showers that develop very rapidly. It is clear though that fluctuations in the propagation and splitting of secondary particles is not a large source of fluctuations in the total shower. This effect is even less important in hadronic showers since the first interaction of the primary

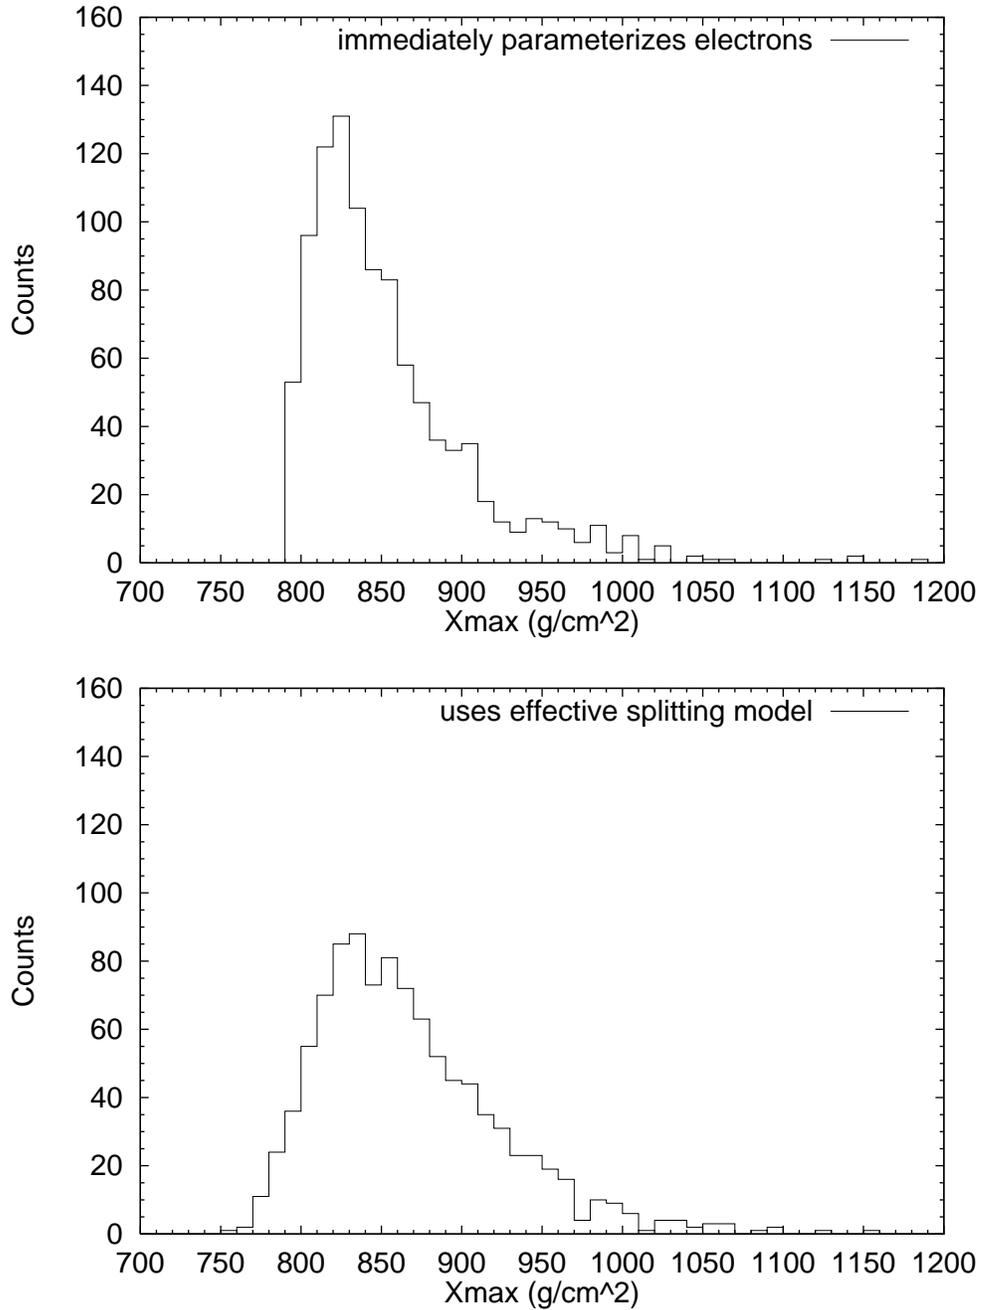


Figure 2.6: Depth at maximum distributions for 10^{18} eV photon initiated showers. In the top histogram electrons are described by their average relative to their depth of creation. The bottom histogram uses the effective splitting model with a simulation threshold of 10^{15} eV. Each set is comprised of 1000 showers.

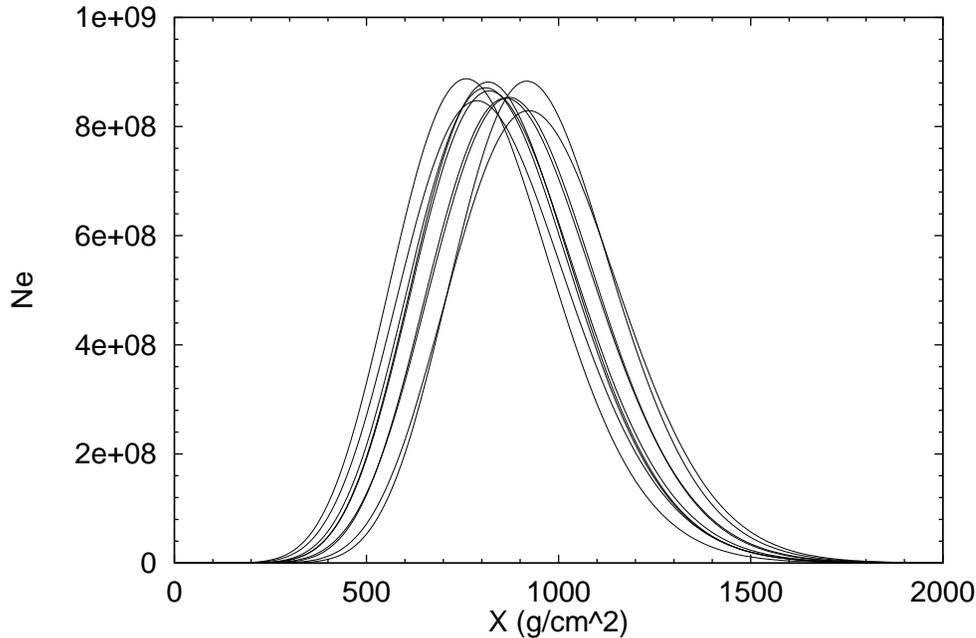


Figure 2.7: Longitudinal profiles for ten 10^{18} eV photon showers. The effective splitting model was used to conduct Monte Carlo simulation down below 10^{15} eV.

can generate hundreds of secondaries (though most of the energy is carried by only a few hadrons).

Figure 2.7 shows the longitudinal profile for ten photon induced showers at 10^{18} eV simulated with the effective splitting model. As discussed earlier, since the length scale of propagation is much smaller than the width of the full shower, fluctuations are characterized by small deviations in the depth at maximum and the size at maximum, but no drastic changes in shape.

2.12 Summary

A modified form of the Greisen formula was used to construct parameterizations for electromagnetic cascades under various initial conditions. Parameterizations were developed by requiring consistency with the prescription $\alpha = 0$ and $\Delta = 0$ for the average development of photon initiated cascades relative to their depth of creation (or injection at the top of the atmosphere). The infrared divergence associated with

Description	α	Δ	error (%)
Photons relative to their depth of creation (the standard Greisen formula).	0	0	—
Neutral Pions relative to their depth of decay into two photons.	0.33	-0.17	0.04
Photons relative to their depth of pair production.	1.10	-0.19	0.05
Electrons relative to their depth of creation	0.80	-0.01	0.03
Effective splitting model with $\lambda_{\text{eff}} = 0.310$ and $\langle x \rangle = 0.116$	0.867	-0.254	0.006

Table 2.2: Summary of the modified Greisen parameters for a variety of initial conditions. There is no error for the first entry since it is an assumption by which the remaining are derived.

bremsstrahlung was handled with an effective splitting model. Optimization yielded a model for which there was near perfect consistency between Monte Carlo simulation and parameterization. Table 2.2 summarizes the results.

Chapter 3

Hadronic Interactions

This chapter develops the tools necessary to simulate interactions of hadrons on nuclei. The emphasis is on developing model which can be easily modified so that various physical assumptions can be explored. A range of models are specified which conservatively bracket the uncertainties associated with features relevant to the highest energy hadronic air showers. A conservative bracketing of uncertainties makes clear what kind of signal indicates radically different physics or exotic types of primaries. While shower development is influenced by only general features of hadronic interactions, information about these features extracted from laboratory experiments is rather scant. This is for a number of reasons:

- The highest center of mass energy achieved thus far in collider experiments, $\sqrt{s} = 900\text{GeV}$, is over two hundred times lower than in an interaction of a primary cosmic ray proton of energy 10^{20}eV in the lab frame with a nucleon from an air nucleus.
- Collider experiments address the central region in rapidity where there is a large number of particles but only a small fraction of the total energy. The bulk of the energy is contained in the high energy particles that are lost down the beam pipe; however, these high energy particles are crucial to air shower development.
- Fixed target experiments that can directly explore the relevant fragmentation region are limited to energies of hundreds of GeV in the lab frame.

- QCD, the underlying theory of hadronic interactions, can not be used to undertake a rigorous calculation of the soft processes since perturbation theory can not be applied due to the large coupling constant. It is these low transverse momentum processes which dominate the high energy interactions in the hadronic core of an air shower.

Still, it is possible to build an approximate range of models based on simple physical ideas which captures all relevant features and is in reasonable accord with experimental data. Arguably, it is actually misleading to build models which fit collider data in detail when the gap that must be bridged to connect with the relevant features to air showers can easily wash away the relevance of such detail. The focus throughout this work is on general trends.

The basic description of a hadronic shower is similar to that of a pure electromagnetic cascade: successive generations of interactions generate an exponentially increasing number of hadrons. A single interaction between a hadron and an air nucleus can generate hundreds of secondaries, the bulk of these being pions. Roughly ten percent of secondaries produced are kaons (production of secondaries with heavier quarks is negligible). Since hadronic interaction properties are similar for all types of mesons, for simplicity the treatment here assumes only pions are generated. Neutral pions at energies below about 10^{19} eV typically decay into two photons, initiating pure electromagnetic subshowers. At energies above about a TeV, charged pions typically interact to generate more pions. Thus the basic picture of a hadronic shower is of a hadronic core which continuously feeds the electromagnetic component through the decay of neutral pions. Unlike pure electromagnetic showers, there is a significant muonic component which is generated by the decay of low energy charged pions, each into a muon and a neutrino. As is demonstrated in the next chapter, at the highest energies, roughly five percent of a primary proton's energy eventually winds up in the form of muons and neutrinos.

3.1 Cross sections

The extrapolation of the proton-air inelastic cross section is the largest source of uncertainty in developing a precise understanding of hadronic shower development at the highest energies. Current models in use have this cross section at 10^{20} eV in the lab frame anywhere in the range from approximately 450 mb to 600 mb.

There exists a number of models which describe accelerator data well where energies probed extend up to a center of mass energy of 1.8 TeV under the assumption that the difference between the $p\bar{p}$ and pp cross sections vanishes at high energies. The highest energy cosmic rays are characterized by a center of mass energy of $\sqrt{s} > 400$ TeV. The Large Hadron Collider (LHC) will take the upper limit to as high as $\sqrt{s} = 14$ TeV. Until then, the only source of information in this energy region is air shower experiments, the interpretations of which are very difficult. The most directly measured quantity, the inelastic proton-air cross section, is characterized by errors of around 100 mb. Analyses have been done independently by Akeno (now part of AGASA) [28] and by Fly's Eye [29]. While the means of extracting the depth at maximum are different for the two types of experiments (ground array and air fluorescence respectively), both focus on an analysis of the tail of the depth at maximum distribution. It is closely related to $\sigma_{p\text{-air}}^{\text{inel}}$ under the assumption that the primary cosmic ray composition contains a significant number of protons. The deeply penetrating tail can be connected with the mean free path of the primary though the nature of development fluctuations does introduce a level of model dependence. The analysis by Fly's Eye restricts $\sigma_{p\text{-air}}^{\text{inel}}$ to 530 ± 100 mb at $\sqrt{s} \approx 30$ TeV [23].

For the modeling of air showers, what is of interest is the inelastic cross section since elastic scattering results in only a small deflection of the projectile and no loss in energy. It is important to describe the inelastic cross sections for protons and pions on air and nucleon targets. While only the cross sections on air are needed to calculate interaction lengths, the ratio between the cross sections on air to that on a single nucleon is important for modeling the hadronization phase of the interaction. The ratio fixes the average number of nucleons from the target which participate in the interaction.

The Glauber multiple scattering formalism connects cross sections on nuclear targets to the underlying cross sections on nucleon targets. The basic picture can be understood with a highly simplified version of the theory [14] where

$$\sigma_{p\text{-air}} = \int d^2b \{1 - \exp[-\sigma_{pp}T(b)]\}. \quad (3.1)$$

The exponential term represents the survival probability for a projectile at impact parameter b where the number density of target nucleons over the path associated with the impact parameter is given by

$$T(b) = \int \rho_N(b) dz. \quad (3.2)$$

Here ρ_N is the number density of nucleons at a distance $r = \sqrt{b^2 + z^2}$ relative to the center of the target nucleus. If $\sigma_{pp}T(b) \ll 1$, the exponential can be expanded giving

$$\sigma_{p\text{-air}} \approx \sigma_{pp} \int d^2b T(b) = A\sigma_{pp} \quad (3.3)$$

where A is the number of nucleons in the target nucleus. In this limit, the probability of interacting with any given nucleon is so small that participation of multiple nucleons in a single interaction is negligible. The configuration of the nucleus is therefore irrelevant. In the other extreme, if $\sigma_{pp}T(b)$ is very large, the exponential can be ignored leaving

$$\sigma_{p\text{-air}} \approx \int_0^{R_A} d^2b \propto A^{2/3}. \quad (3.4)$$

In this limit the target nucleus becomes a perfectly absorbing black disk. The cross section is sensitive only to the geometric cross section that the nucleus presents to the projectile.

A convenient parameterization due to Kopeliovich et al. uses the Glauber formalism and carefully extends it up to high energy [30],

$$\sigma_{p\text{-air}}^{\text{inel}} \approx 507 \text{mb} \left(\frac{\sigma_{pp}^{\text{inel}}}{100 \text{mb}} \right)^{0.529}. \quad (3.5)$$

The relation is nearly model independent over a class of models in which nucleon-nucleon scattering is related to the charge distribution within the nucleon. It is valid for energies greater than about 2 TeV in the lab frame. To deal with low energies

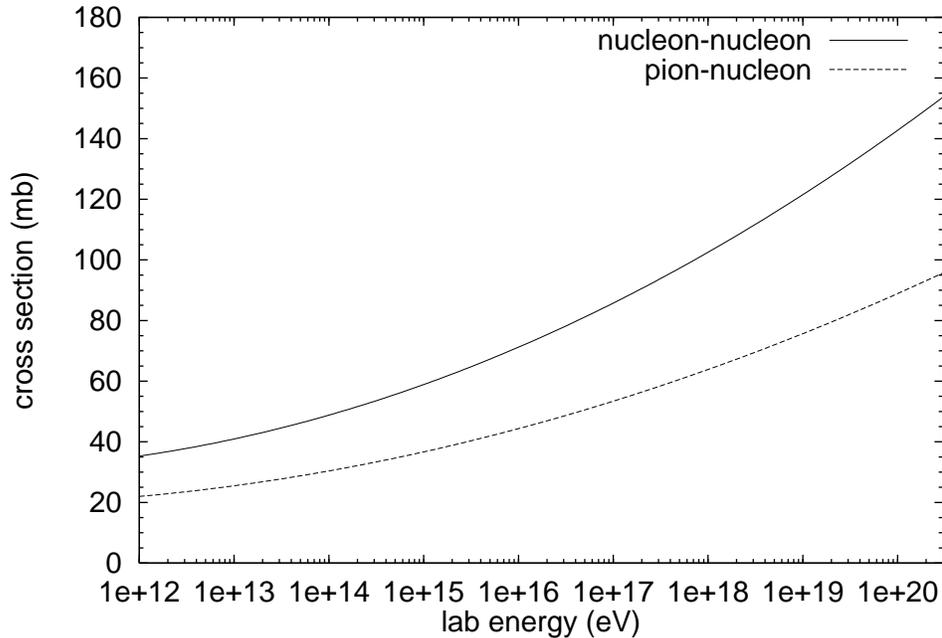


Figure 3.1: Energy dependence of the inelastic cross sections for hadrons on a nucleon target. The ratio of the pion to nucleon cross sections is held at a constant fixed by low energy data.

(small cross sections), the low energy limit of Equation 3.3 is adopted if it gives a smaller cross section than the above parameterization. It is demonstrated in the next chapter that the electromagnetic component of the highest energy showers is totally insensitive to the treatment of the hadronic interactions below about a TeV, and so the crudeness of this extension to low energies is acceptable.

Equation 3.5 reasonably describes any hadronic projectile which is significantly smaller than the target nucleus (the simplified glauher discussion above makes no reference to the structure of the projectile). Consequently, it is adopted to describe pion projectiles as well as nucleon projectiles. This retains the general trend that pions, which are characterized by smaller cross sections, are further away from the black disk limit.

To further simplify the description of hadronic-cross sections, consider a naive parton-picture which decomposes the hadron-nucleon cross section in terms of cross sections describing interactions between valence quarks in the projectile and target.

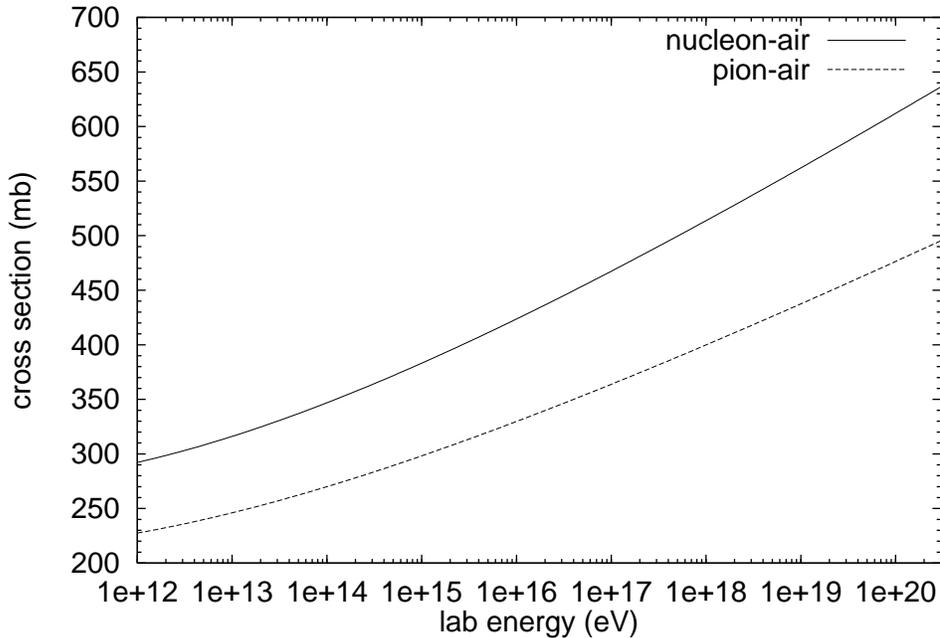


Figure 3.2: Energy dependence of the inelastic cross sections for hadrons on air. Because of the shadowing effects associated with the target nucleus, the ratio between the pion and nucleon cross sections is larger than for on a nucleon target.

This implies

$$\frac{\sigma_{\pi p}}{\sigma_{pp}} \approx \frac{2 \times 3 \times \sigma_{qq}}{3 \times 3 \times \sigma_{qq}} = \frac{2}{3}. \quad (3.6)$$

A simple fit to accelerator data using Regge theory [20] sets the high energy ratio of these cross sections at $13.7/22.0 \approx 0.62$ which is very close to the simple-minded picture above.

Adopting the ratio obtained by the fit to accelerator data and taking into account shadowing effects for a nuclear target with Equation 3.5, all that is needed is the explicit energy dependence of one of the four cross sections from the combinations (pion, nucleon) on (nucleon, air). A possible parameterization in terms of $\ln(s)$ for

the inelastic nucleon-nucleon cross section [31] is ¹,

$$\sigma_{\text{inel}}(s) = 32.4 - 1.2 \ln(s) + 0.21 \ln^2(s), \quad (3.8)$$

where s is in GeV^2 . Figure 3.1 shows the energy dependence of the nucleon and pion cross sections for a nucleon target. Other forms such as a power-law motivated by certain versions of Regge theory describe accelerator data equally well in the region where fits but can extrapolate to quite different cross sections at the highest energies. Consequently, it is important to test the model dependence associated with the choice of extrapolation. Figure 3.2 shows the cross sections on air.

So that cross sections can be easily modified to test model dependence, the following form for the nucleon-nucleon inelastic cross section is adopted:

$$\sigma_{\text{inel}}(s) = 30.7 + ax^2\Theta(x), \quad (3.9)$$

where

$$x = \ln(s) - 2.86 \quad (3.10)$$

with s in units of GeV^2 . With $a = 0.21$, this is the same as Equation 3.8 but normalized so that $x = 0$ corresponds to the minimum. At energies below the minimum the Θ function holds the cross section constant. This allows different choices for a while still preserving the same low energy limit and a smooth evolution from $x = 0$. It is shown in the next chapter that shower development at the highest energies is unaffected by how the physics is modeled up to about a TeV: all particles below this energy can be ignored with negligible effect on the longitudinal profile.

The cross sections on air are translated to an interaction length through the relation,

$$\lambda = \frac{2.4 \times 10^4}{\sigma_{\text{air}}}, \quad (3.11)$$

¹Collider experiments typically cast observations in terms of the variable s which is the center of mass energy squared. At high energies compared to the proton mass (assuming a proton target),

$$s \approx 2m_p E_{\text{lab}}. \quad (3.7)$$

where $m_p \approx 0.938 \text{ GeV}$. Since air shower experiments take place in the lab frame (with the atmosphere at rest), when plotting the energy dependence of any quantity, the convention adopted in this work is to plot everything in terms of E_{lab} over the range from a TeV up to the largest energy yet detected, $3 \times 10^{20} \text{ eV}$.

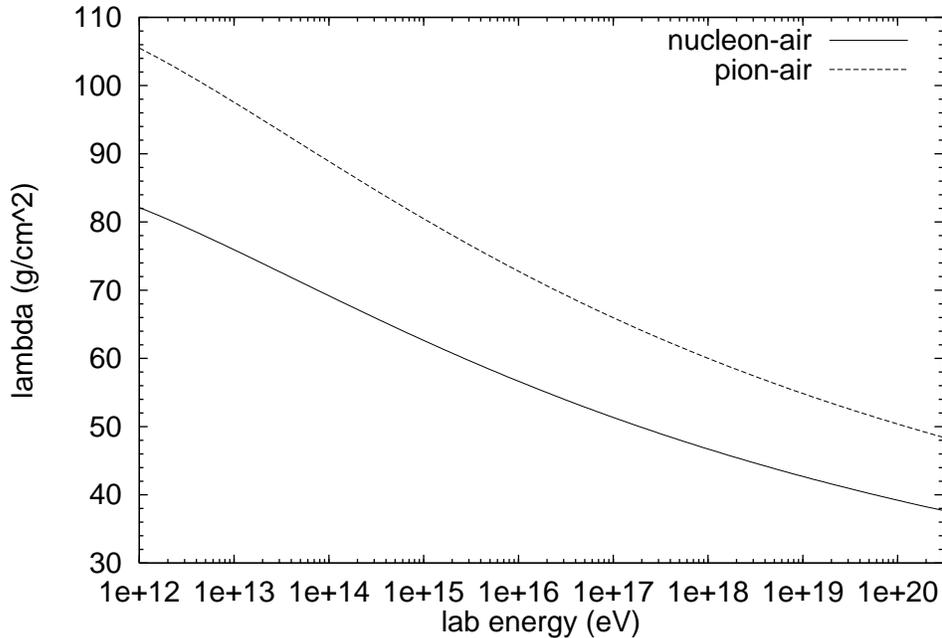


Figure 3.3: Energy dependence of the interaction length for hadrons on air.

where the cross section is in mb and the interaction length is in g/cm^2 . This quantity is shown Figure 3.3. It is this length that sets the scale for fluctuations in shower development since it is the depth of the first interaction that represents the dominant source of fluctuations in shower development.

3.2 Hadronization: general considerations

The building of an event generator along with a model of the underlying physics is largely phenomenological. Of course it is desirable that a model implies results that are in accord with experimental data at a level suitable for air shower simulation. It is also worthwhile to ask if the model extrapolates to the highest energies in a meaningful manner. The current state of non-perturbative QCD does not allow for any precise calculation for the processes relevant to air shower development. At best, a basically correct model will at some point need to make use of an ansatz which is based on ideas that are qualitative. At worst, an event generator may fit accelerator data

precisely but may be meaningless when extrapolated to higher energies or applied to the forward fragmentation region. Since there are a number of important features which must be incorporated, it is perhaps best to begin with the description of a simple event generator followed by a discussion of how it is related to the underlying physics and how it can be modified to best model the features relevant to air shower development.

3.2.1 The Hillas splitting algorithm

The splitting algorithm due to Hillas [32] is an example of an efficient technique which is easily adapted to explore physical assumptions and their influence on air showers. A preliminary description of nucleon-nucleon interaction carried out in the lab frame (a projectile on a fixed target) is:

1. Randomly split the available energy into two pieces using a uniform deviate to assign a fractional energy.
2. Assign one piece as the energy of the leading nucleon.
3. Split the remaining piece into 2^N fragments, with $N = 2$, by conducting N presplittings by which pieces are successively split in two with uniform deviates.
4. Each of the 2^N fragments undergoes an iterative process by which the fragment energy is split using a uniform deviate. One piece is assigned as a pion and the remaining piece is assigned as a fragment which continues to produce pions until the remaining energy is less than the pion mass.

This was the original form used in early simulations conducted by Hillas in the modeling of extensive air showers. An important technical aspect of the algorithm is that it automatically conserves energy to the level required for the task at hand. This type of splitting algorithm does permit the generation of packets of energy less than the pion mass. This is dealt with by ignoring such packets. Tracking the total energy ignored in this manner allows an assessment of the error incurred. It is shown in the next chapter that the error is negligible at the energies of interest. Another technical

aspect worth mentioning is that algorithms of this type naturally generate the highest energy particles in the earliest stages. This allow for efficient simulation when the lowest energy of interest may be much higher than the pion mass: all fragments which drop below the lowest energy of interest can be ignored.

There are three main features related to the underlying physics to take note of at this point:

1. A single particle is generated which on average carries an energy equal to that of all the remaining particles combined.
2. The number density of particles diverges as the fractional energy relative to the parent projectile approaches zero.
3. While the average multiplicity (total number of particles) grows with the energy of the initial nucleon, the x_{lab} (fractional energy relative to the projectile energy) distribution does not change with energy except for a threshold effect. In other words, the interaction scales with energy except near $x_{\text{lab}} = m_{\pi}/E_0$.

These three aspects parallel the composite structure of the projectile nucleon. A nucleon is composed of three valence quarks which carry roughly half of the nucleon's energy. The remainder of the energy is contained in a sea of particles composed of gluons and quark-antiquark pairs. Collectively these point-like particles are often referred to as partons. Qualitatively it is expected that the momentum distribution of generated particles reflects the momentum distribution of partons in the projectile nucleon and to some extent the target as well [14]. The flat distribution of the leading particle reflects the valence component of the projectile while the numerous secondaries reflect the sea for which the distribution of partons has a small x behavior which is divergent.

3.2.2 Analytic description of fragments

The fact that all splittings are carried out using a flat distribution allows for analytic calculations of momentum distributions and multiplicities. Focusing attention

on a single fragment which is undergoing $fragment \rightarrow fragment + pion$, the probability distribution for the fractional energy of the fragment emerging from a single splitting is the same as that for the pion. Consequently, the probability of generating on the n th splitting a pion with fractional energy in the range x to $x + dx$ is given by

$$p_n(x)dx = \int_x^1 p_{n-1}(x') \frac{dx'}{x'}. \quad (3.12)$$

Here x is relative to the original fragment energy before any splittings. The integral is a convolution over all possible fragments from the previous splitting which have sufficient energy to produce a pion with fractional energy x . The factor $\frac{dx'}{x'}$ is the probability of producing a pion in the requisite range given a fragment that is capable of producing such a pion. The symmetry between fragments and pions allows $p_{n-1}(x)$ to be viewed as a fragment probability density while p_n is viewed as a pion probability density. With the constraint that $p_1(x) = 1$, the probability density at any splitting is calculated to be,

$$p_n(x)dx = \frac{\left(\ln \frac{1}{x}\right)^{n-1}}{(n-1)!}. \quad (3.13)$$

The x distribution of pions associated with a single fragment is then obtained by summing over all splittings,

$$\frac{dn_\pi}{dx} = \sum_{n=1}^{\infty} p_n(x) = \sum_{m=0}^{\infty} \frac{\left(\ln \frac{1}{x}\right)^m}{m!} = \frac{1}{x}. \quad (3.14)$$

The average multiplicity of particles above a threshold x_{\min} is then given by,

$$\langle n_\pi \rangle = \int_{x_{\min}}^1 \frac{dn_\pi}{dx} dx = \ln \frac{1}{x_{\min}}. \quad (3.15)$$

The effect of incorporating presplittings has not been taken into account. However, regardless of how the set of fragments are generated it is clear that multiplicities grow logarithmically with energy as long as the algorithm which generates the fragments scales with energy. In cases where the algorithm may exhibit scaling violations there can be deviations from strict logarithmic growth, but nonetheless the logarithm of the energy still sets the scale for multiplicities.

A convenient approach to developing models of hadronic interactions is to express event generation in terms of fragments which hadronize as described above. Equations 3.14 and 3.15 can be useful tools by convoluting these results with probability distributions which describe the generation of fragments for any particular model.

3.2.3 String fragmentation

While the distribution of hadrons from an interaction does seem to reflect the composite structure of the projectile, there is a related but distinct viewpoint which interprets the distribution in terms of the fragmentation of QCD strings. The picture which is briefly discussed here is the Lund algorithm [33] which has been used extensively to describe jets in accelerator experiments. For the purposes of illustration, the hadronization associated with a back-to-back quark-antiquark jet pair as produced in electron-positron annihilation is a particularly clean situation in which a string model can be employed. The quark and antiquark move apart with a string representing the linear confinement potential of QCD stretched between them. As they move apart, the potential energy stored in the string increases eventually causing a break which results in the creation of a new quark-antiquark pair. Provided the invariant masses of the resulting systems are large enough, additional breaks can form. Hadrons are formed by the grouping of a quark from one break with the antiquark from a neighboring break. The method of determining break points in the string is guided by demanding that the method be Lorentz covariant and that the final hadrons emerge with their correct masses. The formation of break points can not be understood in terms of any causally connected progression in time since the breaks have space-like separations between them.

The precise formulation of string fragmentation is cumbersome and inefficient for Monte Carlo implementation. An approximate iterative scheme which produces results experimentally indistinguishable from the more rigorous formulation is usually implemented. This iterative scheme works much like the hadronization of fragments above: $string \rightarrow string + hadron$. The main difference is that the distribution

employed is not flat but takes the form,

$$p(x)dx \propto \frac{(1-x)^a}{x} \exp(-bm_T^2/x), \quad (3.16)$$

where a and b are free parameters to be fit by experiment and m_T is the transverse mass ($m_T^2 = E^2 - p_z^2$, where the jet axis is the z-axis), which is sampled from an appropriate distribution. Since only one-dimensional shower generation is being considered, the effect of a varying transverse mass is inconsequential, all transverse masses are of the order of hundreds of MeV. Technically what is sampled is not the fractional energy but a quantity which guarantees the Lorentz covariance of the procedure (the formulation makes use of light-cone variables). However, in the limit where the energy remaining is large compared to the hadron masses, x can be interpreted as the fractional energy. The form of equation 3.16 is fixed by requiring the same results on average whether fragmentation proceeds by starting from the quark end or by starting from the anti-quark end. In other words, fragmentation is left-right symmetric.

The precise shape of Equation 3.16 is of little concern here. What is important is that the process of string fragmentation can be described by an iterative procedure of the same type as the Hillas splitting algorithm. Regardless of the shape of the probability distribution, the overall distribution associated with a fully fragmented string looks essentially the same as that achieved with any probability distribution. It is important to realize that the process of iteratively generating hadrons is of no physical relevance since the break points in the string are causally disconnected: all that is relevant is the final distribution. The suppression of high energy hadrons is controlled in the Lund algorithm by the shape of the probability distribution; Hillas' splitting algorithm achieves the same end by conducting presplittings. Arguably string fragmentation is an algorithm more motivated by the general structure of QCD; however, it still requires modifications which allow for a leading particle when applied to soft collisions. In a context where there is no concern with maintaining the correlation between longitudinal and transverse degrees of freedom in hadronization, it is not clear to what degree soft collisions parallel the hadronization of hard jets. The confining nature of QCD does suggest that there is a parallel at some level, however it is probably best to regard the procedure as a motivated ansatz.

In summary, whether hadronization is physically interpreted as a process which mirrors the distribution of partons inside the composite hadronic projectile or as a process related to QCD string fragmentation, the final end is the same: both are characterized by a number density which diverges exponentially as $x_{\text{lab}} \rightarrow 0$. This divergence is most easily characterized using an iterative procedure by which the projectile energy is decomposed into fragments which undergo $\text{fragment} \rightarrow \text{fragment} + \text{pion}$. The use of presplittings is a convenient tool which achieves the suppression of high energy pions while still using a flat distribution to conduct all splittings.

3.3 The leading particle

The use of a flat distribution to characterize the leading nucleon is supported experimentally in fixed target experiments extending up to 400 GeV in the lab which see a flat nucleon distribution in the forward fragmentation region in pp collisions [34, 35]. There is a slight turn up at low values of x_{lab} which is presumably associated with the generation of nucleon-antinucleon pairs. At high values of x_{lab} the data suggests that the flat distribution tapers off though it is difficult to interpret the distribution above $x_{\text{lab}} \approx 0.85$ due to triggering conditions. With regard to pion projectiles, the distribution of the highest energy pions is softer though the notion of a leading particle is vague experimentally since the type of the projectile is the same as the secondaries. Still the distribution of pions can be examined with particular focus on high values of x_{lab} . Consequently, it is desirable to modify the basic splitting algorithm to generate a softer leading particle distribution for pion projectiles and possibly a diminished probability for high values of x_{lab} for proton projectiles.

A modification similar to that used in the Sibyll event generator [36] involves a simple consideration of the valence component of nucleons versus pions. It is natural to assume that some fraction of the partons in the projectile are spectators in the interaction. These spectators are presumably associated with the leading particle. Assume that in some sense only one of the valence quarks participates in the interaction. Rather than sample a flat distribution from the total energy of the projectile, only sample from the energy of the spectators. This can be accomplished by first

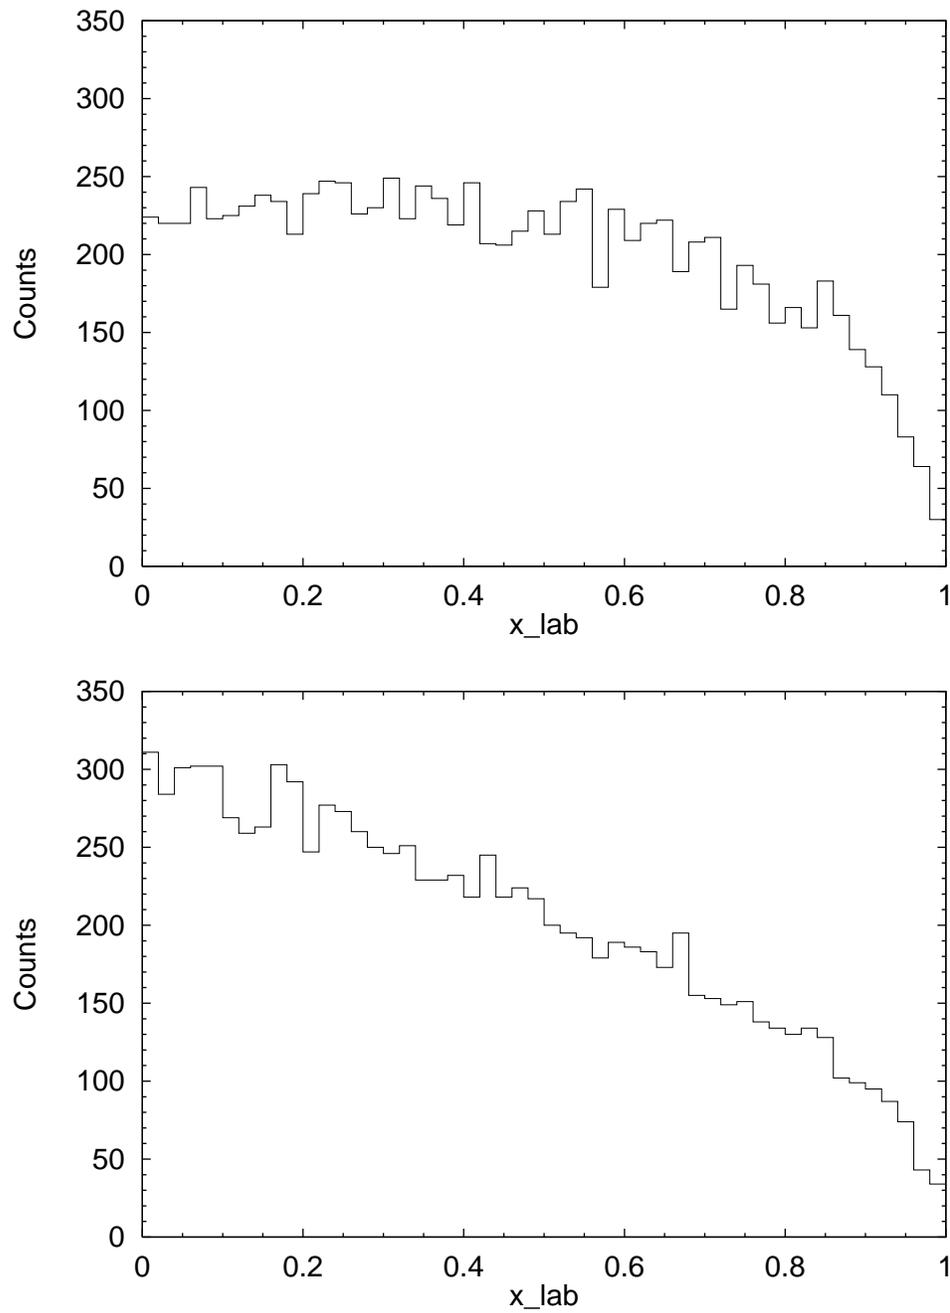


Figure 3.4: Leading particle distributions for nucleon projectiles on a nucleon target (top) and pion projectiles on a nucleon target (bottom). Each set is comprised of 10,000 simulated interactions.

sampling a valence quark fragment from the projectile according to the distribution,

$$p(x)dx \propto \frac{(1-x)^{2n_s-1}}{\sqrt{x}}, \quad (3.17)$$

where n_s is the number of valence quarks which are spectators. The numerator suppresses the probability that a single valence quark carries a high value of x due to the fact that there are other valence quarks which carry significant fractions of the projectile's energy. Counting rules have been proposed which suggest that this numerator expresses the high x dependence of quark momentum distributions [37]. The denominator gives the expected low x behavior for valence quarks. The energy sampled in this manner is then deducted from the energy store after which the leading particle is generated as normal. Following the generation of the leading particle, all remaining energy is lumped into the energy store which then undergoes presplittings. This is a simple ansatz which leads to a slightly softer distribution for pions while retaining a relatively flat distribution for leading nucleons at moderate and low values of x .

Often hadronization models are characterized by the inelasticity which is defined as the fractional energy given up by the leading particle. In the original Hillas algorithm, the inelasticity is 0.50 while in this modified algorithm it can be expressed as,

$$K = 1.0 - \int_0^1 0.5(1.0 - x)p(x)dx. \quad (3.18)$$

This gives 0.6 for pions and approximately 0.56 for nucleons. The inelasticity is a fundamental quantity associated with air shower development since it is one of the main factors which controls the rate at which energy is channelled from the hadronic core into the electromagnetic cascade. Figure 3.4 shows leading particle distributions for nucleons and pions. As desired, the leading nucleon distribution is flat up until large values of x_{lab} and the pion distribution is significantly softer.

3.4 The distribution of secondaries

Along with interaction lengths and inelasticities, the energy distribution of secondaries in the forward fragmentation region is fundamental in controlling the devel-

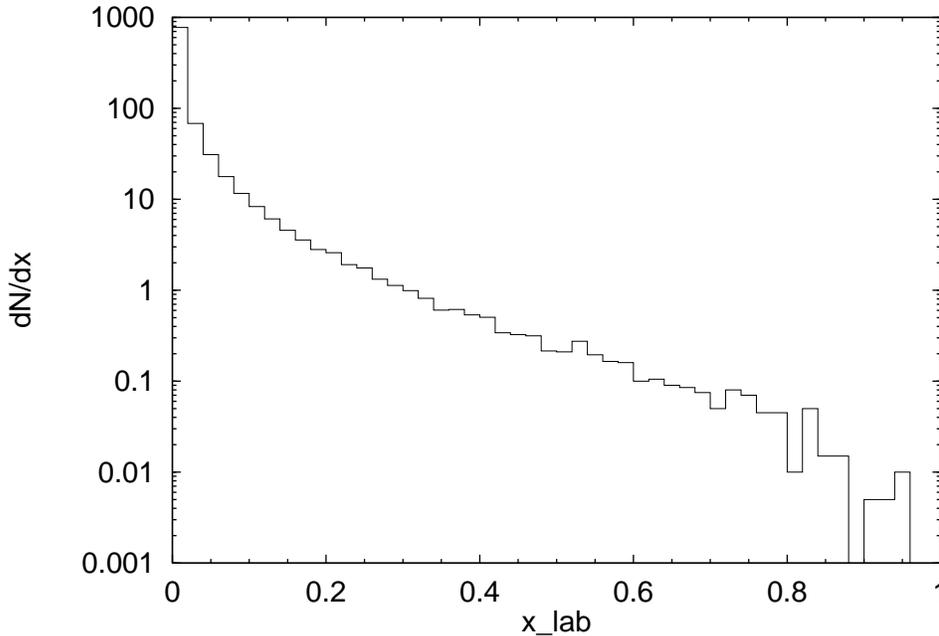


Figure 3.5: x_{lab} distribution of charged pions for nucleon projectiles on a nucleon target. The histogram is constructed from 10,000 simulated interactions.

oment of the hadronic core. A related quantity is the total multiplicity of particles. However, it is important to note that multiplicities are important because of what they imply for the distribution of secondaries with large values of x_{lab} . This connection between multiplicities and the fragmentation region is very model dependent. Given no direct way of experimentally assessing this connection and the current inability to rigorously establish it on theoretical grounds, it is important not to place too much emphasis on a precise fitting of particle multiplicities.

While the original statement of Hillas' splitting algorithm does a reasonably good job of characterizing secondaries, it can be improved upon in three major ways:

1. Fluctuations in total multiplicities are underestimated.
2. There is a mild discrepancy in x_{lab} distributions with measurements from fixed target experiments which extend up to 400 GeV in the lab.
3. Above about $\sqrt{s} \approx 100$ GeV, collider experiments show that multiplicities grow

<i>type</i>	$\langle N \rangle$	<i>error</i>
π^0	1.94	0.06
π^+	2.05	0.055
π^-	1.67	0.04
K^+	0.166	0.08
K^-	0.112	0.055
p	0.60	0.097
\bar{p}	0.031	0.001

Table 3.1: Average multiplicity of particles observed in the forward fragmentation region as observed by the LEBC-EHS collaboration for pp interactions at 400 GeV in the lab frame.

faster than linearly with the logarithm of the energy though the log of the energy still sets the appropriate scale.

The underestimation of low multiplicity events might have effects on the deeply penetrating tail of the depth at maximum distribution for proton showers. In general, interaction generators in use have trouble modeling low multiplicity events. Larger fluctuations can be incorporated by sampling the number of presplittings from a distribution rather than keeping the number fixed at two as in the original statement of Hillas' splitting algorithm. This is achieved by sampling the number of presplittings from a Poisson distribution. This leaves the average number of presplittings as a free parameter.

This parameter is fixed by fitting the shape of the x_{lab} distribution of charged pions in pp interactions which have been examined in fixed target experiments up to 400 GeV [34, 35]. A convenient presentation of the experimentally measured charged pion distribution appears in Fletcher [35]. In order to make a comparison, only a fraction of the mesons generated are assigned as charged pions when conducting simulations for a trial model. The fraction of mesons assigned as charged pions is extracted from the average multiplicities of various types of particles [35]. These are indicated in Table 3.1. The forward fragmentation region corresponds to particles with a positive longitudinal momentum when viewed in the center of mass frame. The total multiplicity is larger by a factor of two to account for the fragmentation of

the target. The asymmetry between positive and negative pions can be understood in terms of the charge of the projectile and the emerging leading nucleon. While the influence of kaons and proton-antiproton pairs is included in the charged pion assignment here to allow a close comparison, the final model generates only charged and neutral pions with the two to one ratio (all types equally probable) suggested by observation. The influence of kaons is small since their properties are similar to those of pions and their production probability is an order of magnitude less. The modeling of such effects are not of much concern here since the focus is on overall trends. Also be aware that the figures appearing in Fletcher plot x_F which corresponds to the fractional energies in the forward region of the center of mass frame. This quantity is related to the fractional energy as it appears in the lab by

$$x_{\text{lab}} = \frac{1}{2} \left(\sqrt{x_F^2 + \frac{4m_{\text{T}}^2}{s}} + x_F \right) \approx x_F, \quad (3.19)$$

where $m_{\text{T}} = \sqrt{m^2 + p_{\text{T}}^2}$ is the transverse mass. The approximation is valid when $x_F \gg 4m_{\text{T}}^2/s$. Since transverse masses are of the order of hundreds of MeV, there is an appreciable difference between these quantities in only the first few bins; these bins should not be used in the comparison. Excellent agreement with experiment is obtained with a value of 2.3 ± 0.1 for the average number of presplitting. As expected, it is close to the fixed number of presplittings in the original splitting algorithm.

If the average number of presplittings is taken to be independent of energy, by Equation 3.15 the multiplicity grows logarithmically with the energy. Departures from this behavior necessarily involve scaling violations of some sort.

3.5 Scaling violations

It is difficult to directly fit multiplicities since the adopted splitting algorithm is not expected to accurately describe the fragmentation of the target. These low values of x_{lab} are not significant in any case since this region contributes negligibly to the development of an air shower. Consequently, while it is important that the algorithm allows for low multiplicity events, the consideration of multiplicities should

be concerned with general trends rather than precise fitting. The observation of mild scaling violations around $\sqrt{s} \approx 100$ GeV and above is important in so far as it carries implications for the forward fragmentation region which accounts for only a small number of particles. However, it is far from clear how to establish the connection between the central and fragmentation regions. Here, two extreme viewpoints which bracket an accurate description of the physics are examined. Applied to air shower simulations, the relevance of the uncertainties can be assessed.

The first extreme viewpoint assumes that the scaling violations observed in the central region can, to a good approximation, be ignored in the context of hadronic showers. This is an extreme characterization of the case where the high energy particles emerging from an interaction are unaffected by the participating portion of the projectile which accounts for large multiplicities but only a small fraction of the projectile's energy. With such an approximation, the hadronization algorithm is implemented in the same manner regardless of the projectile energy. An implication of this is that the inelasticity is approximately constant as a function of energy. This roughly parallels the Sibyll event generator which gives only a mild rise in the inelasticity due to minijets resulting from the hard scattering of low x_F partons. In minijet models [38], as the interaction energy is raised, minijet production increases but the scattering partons are characterized by diminishing values of x_F due to the QCD evolution of structure functions. The basic idea is that there are parton-parton scatterings which are hard enough to validate the use of perturbation theory. An increase in the energy resolves lower x_F partons, increasing the number available for hard scattering. The UA1 experiment does see evidence for the production of minijets at the level of about 10% of the inelastic cross section [39], suggesting that this model is at least qualitatively correct. The use of a limiting scaling model places a strict lower bound on the effects on the forward fragmentation region.

The second extreme viewpoint assumes there is a strong correlation between the central region and the fragmentation region. This can be accomplished by correlating the inelasticity with the sampling of presplittings. Physically, the strongest correlation which can be reasonably imposed is to assume that the highest energy pions generated in an interaction remain close in energy to the leading particle as the energy is raised.

One way of imposing this is to take the average energies of a first rank pion (the first pion created from a fragment which has just emerged from the presplitting phase) and the leading particle to have a ratio independent of energy. Since each presplitting on average splits the energy in half, this relationship can be expressed as,

$$\left(\frac{1}{2}\right)^N K \approx C(1 - K), \quad (3.20)$$

where N is the average number of presplittings, K is the inelasticity, and C is fixed by the low energy model of pp interactions. The analogous model in this extreme case is the KNP model [30] which uses a naive multiple scattering model to relate the inelasticity to the inelastic cross section. The result gives a rapidly increasing inelasticity reaching approximately 0.8 at 10^{18} eV for pp collisions in the laboratory frame.

The algorithm which generates the leading particle must be modified in a manner which takes into account the rise in the number of presplittings at high energies while gradually approaching the scaling algorithm as the energy is reduced. This is accomplished by sampling the distribution

$$p(x)dx \propto (1 - x)^\alpha, \quad (3.21)$$

where x is a fractional energy which is removed from the available energy before invoking the standard procedure for leading particle generation. The presplittings are then conducted on the initial projectile energy minus the leading particle energy. By calculating the average of $(1 - x)$ for the above distribution, α is fixed according to

$$1 - \bar{x} = \frac{\alpha + 1}{\alpha + 2}, \quad (3.22)$$

where \bar{x} satisfies the relation,

$$K = (1 - \bar{x})K_0 + \bar{x} \quad (3.23)$$

and K_0 is the scaling value for the inelasticity. All that remains is to fix the energy dependence of the average number of presplittings.

While precise fitting of multiplicities can not be done meaningfully due to the crudeness of the hadronization algorithm, enough information can be extracted from

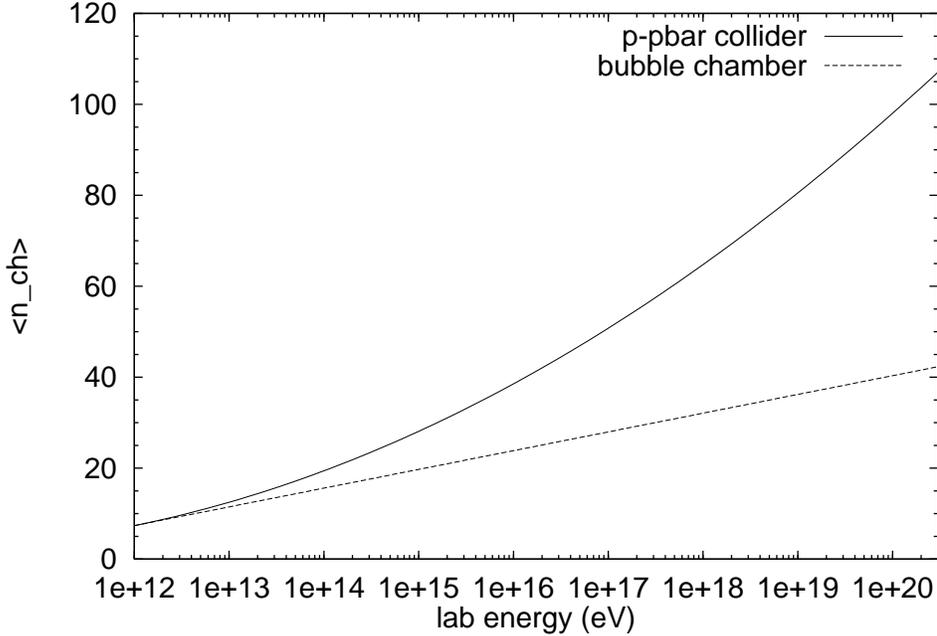


Figure 3.6: Two fits to the average charged multiplicity. The fit which grows logarithmically agrees well with bubble chamber data which extends up to about 5×10^{11} eV in the lab. The fit which grows as the log squared takes into account collider data which extends up to energies corresponding to about 4×10^{14} eV in the lab.

the measured energy dependence of multiplicities to fix the energy dependence of the average number of presplittings. The resulting model extrapolated to the highest energies probably overestimates the extent to which scaling violations carry over to the fragmentation region, but nonetheless it provides a reasonable upper bound.

The average multiplicity of charged particles as observed in bubble chamber experiments which extend up to $\sqrt{s} = 30$ GeV appears to grow as $\ln s$:

$$\langle n_{ch} \rangle = (-2.9 \pm 0.3) + (1.79 \pm 0.5) \ln s, \quad (3.24)$$

where s is measured in units of GeV^2 [40]. Data from collider experiments, ISR and UA5 [41], which examine $p\bar{p}$ interactions extending up to $\sqrt{s} = 900$ GeV observe a $(\ln s)^2$ growth:

$$\langle n_{ch} \rangle = (2.7 \pm 0.7) + (-0.03 \pm 0.21) \ln s + (0.167 \pm 0.016)(\ln s)^2. \quad (3.25)$$

The ratio between these two fits gives information about the degree of scaling vio-

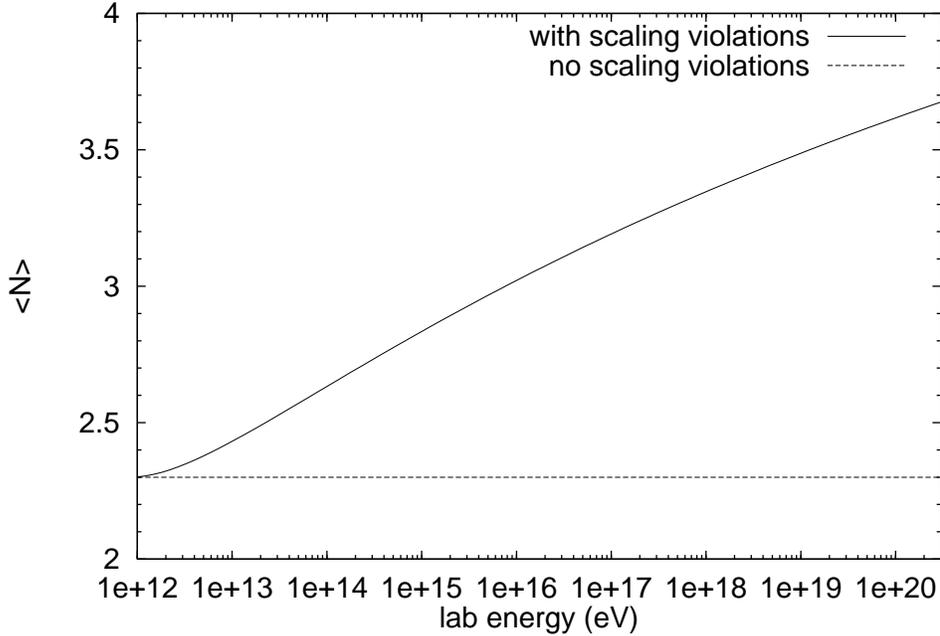


Figure 3.7: The average number of presplittings for the description of pp interactions in the limiting model where the forward fragmentation region is strongly correlated with the central region.

lations. The experimental fits need to be normalized such that they agree at some low energy value. This has been done by adjusting the constant factors in the fits in proportion to their respective error bars such that the fits agree at the energy for which they share the same slope. This occurs at approximately 840 GeV in the lab and requires an increase of about 0.19 in the $\ln s$ fit and a reduction of about 0.45 in the $(\ln s)^2$ fit. These corrections are well within quoted errors. Below this energy where the slopes agree, it is assumed that the scaling model adequately describes the physics. The two resulting fits are shown in Figure 3.6 which extrapolates the fits up to the energies relevant to the highest energy air showers.

Using Equation 3.15 along with an estimate of the average number of fragments and their average energy, an approximate expression for the pion multiplicity is given by

$$\langle n_\pi \rangle \approx 2^N \ln \left(\frac{KE}{2^N m_\pi} \right), \quad (3.26)$$

where N is the average number of presplittings. The ratio between the fit to the

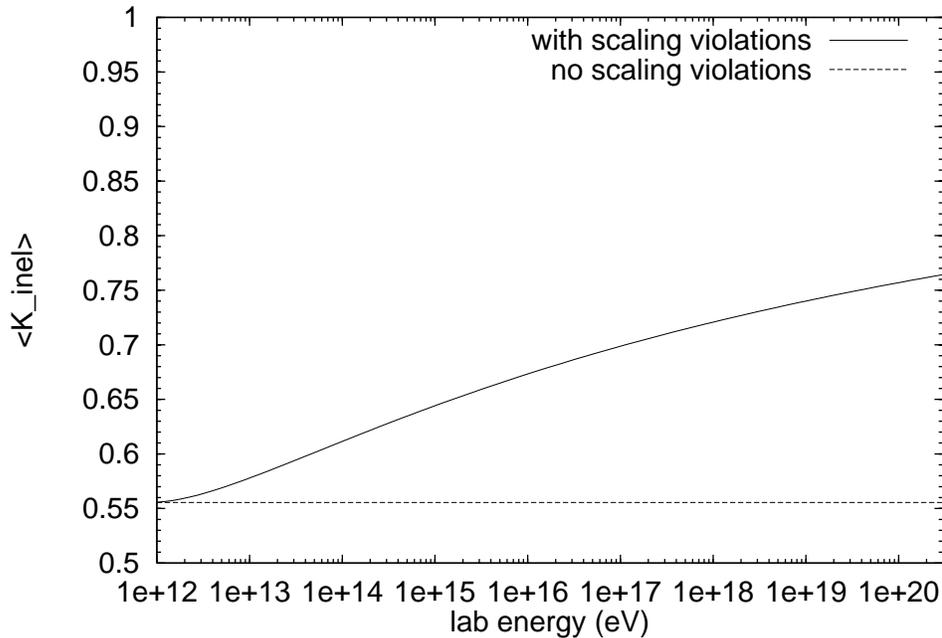


Figure 3.8: Energy dependence of the inelasticity in pp interactions in the limiting model where the forward fragmentation region is strongly correlated with the central region.

charged multiplicity which includes collider data to the fit of just the bubble chamber data is then given by,

$$R \approx 2^{N-N_0} \ln \left(\frac{KE}{2^N m_\pi} \right) / \ln \left(\frac{K_0 E}{2^{N_0} m_\pi} \right), \quad (3.27)$$

where the 0 subscript denotes energy independent scaling values. The ratio, R , is fixed independently by the experimental fits. Equation 3.20 can be used to eliminate K from Equation 3.27 which then allows N , the average number of presplittings, to be solved numerically. Solutions have been tabulated at a number of discrete energies at a level which allows for accurate interpolation. Results are shown in Figure 3.7 which gives the average number of presplittings as a function of energy, and in Figure 3.8 which shows the analogous plot for the inelasticity.

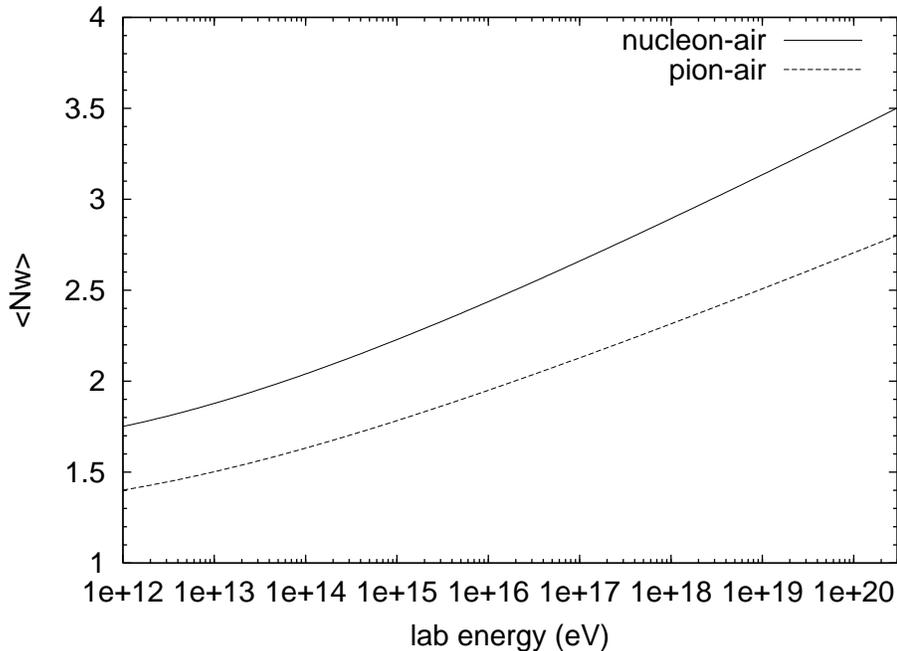


Figure 3.9: Energy dependence of the average number of wounded target nucleons, those that participate inelastically, for hadrons on air.

3.6 Nuclear target effects

The treatment of nucleon-nucleon interactions must be extended to include nuclear target effects. The dominant component of nuclei found in air is nitrogen, and so whenever a fixed atomic mass is needed, $A = 14$ is used. In expressions which permit an average atomic mass, a value of 14.5 is adopted. While the effect on the total inelastic cross sections has been discussed above, it remains to be seen what this implies for hadronization. A multiple scattering formalism [42, 43] which treats nucleon-nucleus interactions in terms of the underlying nucleon-nucleon interactions gives a simple geometric formula for the average number of wounded nucleons,

$$N_w = A\sigma_{pp}^{\text{inel}}/\sigma_{p-\text{air}}^{\text{inel}}. \quad (3.28)$$

A wounded nucleon is a nucleon from the target which participates inelastically. Converting this relation into a hadronization scheme for a single interaction requires two basic steps:

1. Sample a discrete number of participating target nucleons in a single interaction such that the correct average is obtained over many such interactions.
2. Calculate the effect each wounded nucleon has on the hadronization of the projectile.

3.6.1 Sampling the number of wounded nucleons

Rather than embark upon a detailed implementation of the multiple scattering formalism, a simple scheme which captures general trends is adopted. Besides correctly giving the average number of participating nucleons over many interactions, the sampling scheme incorporates the fact that most collisions are only glancing giving rise to inelastic interactions with usually only a single nucleon on the periphery of the target nucleus. The probability that an interaction is characterized by an impact parameter b is characterized roughly by the relation

$$p(b)db \propto \Theta(R_A - b)2\pi b db, \quad (3.29)$$

where R_A is an effective radius for the nucleus. This geometric picture models the nucleus as an absorbing black disk where intersection with the nucleus always yields an interaction. The probability density is largest at R_A since this value for the impact parameter accounts for the largest fraction of the cross sectional area which the target nucleus presents to the projectile. At large impact parameters, the projectile trajectory only passes through a small number of target nucleons.

A simple way of emphasizing peripheral collisions samples an intermediate average number of wounded nucleons corresponding roughly to that for a fixed impact parameter. Actually, one less than the average is sampled since an interaction implies at least one wounded nucleon. The manner of sampling the intermediate average can be tailored to emphasize low numbers of additional participants while still maintaining the correct overall average. This is achieved by sampling the probability distribution,

$$p(n)dn \propto \left(1 - \frac{n}{n_{\max}}\right)^\alpha, \quad (3.30)$$

where $n_{\max} = 13$ for nitrogen, and α is fixed to give the correct average. For any reasonable extrapolation of hadronic cross sections, $\alpha > 0$, and so this distribution always peaks at $n = 0$ with the probability falling to zero at $n = n_{\max}$. From this intermediate average, a binomial distribution with maximum n_{\max} is sampled. Adding one to the resulting number gives the total number of wounded nucleons in the single interaction.

3.6.2 Hadronization effects

It is unclear how the additional participating target nucleons effect the hadronization of the projectile. A picture which is often adopted for the sake of simplicity is the so-called leading-particle cascade [30, 44] where each of the wounded nucleons generates a full nucleon-nucleon interaction with the leading particle always serving as the projectile. Extrapolated to energies where there are significant numbers of wounded nucleons, the inelasticity associated with the p -air interaction drops as the interaction energy is raised, but the x_{lab} distribution for pions actually increases for large values of x_{lab} . This is because pions resulting from each wounded-nucleon interaction do not participate any further. Consequently, the pp secondary distribution of pions is generated with the first wounded nucleon with additional contributions from the remaining wounded nucleons. Such a picture connects the inelasticity with the secondary distribution of pions in an unnatural way. Furthermore, low energy data on proton-nucleus interactions clearly indicates that the inelasticities obtained with the leading-particle cascade are untenable. This rules out statistical models [45] since they require a leading particle cascade so as not to imply showers which penetrate deeper than iron [44]. The leading particle cascade counteracts the decrease of the inelasticity in hadron-nucleon interactions (a decreasing inelasticity is a general feature of statistical models). Therefore, despite the fact that there are a number of models which adopt this picture, it is not investigated here any further even as a limiting model.

As with caling violations in pp interactions, nuclear target effects are investigated in two limits. The first limit considers the effects of a nuclear target to be negligible on

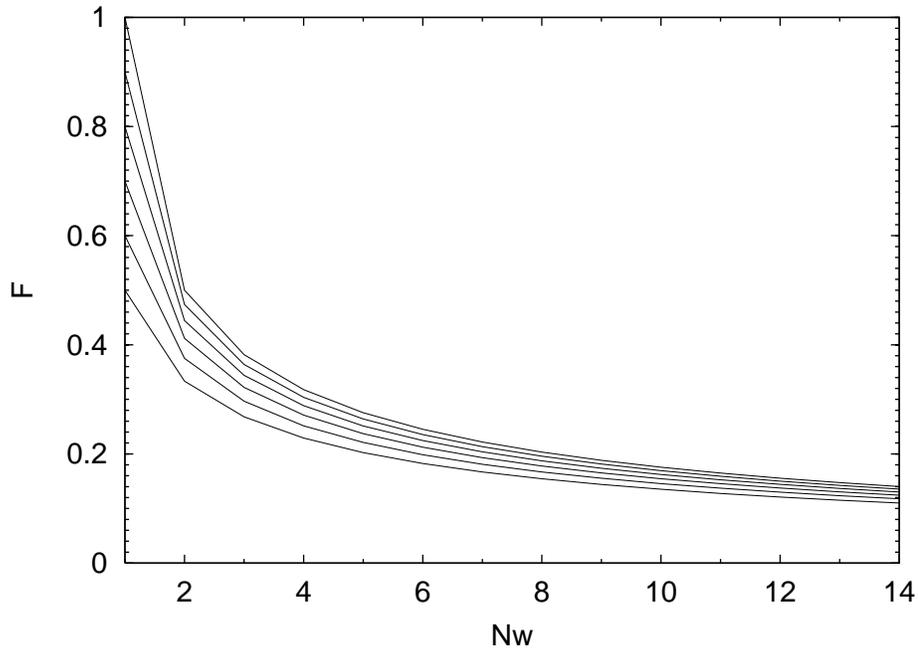


Figure 3.10: The fractional energy for presplits as a function of the number of wounded nucleons. The value at $N_w = 1$ gives the inelasticity. The lines are drawn for visual clarity, only discrete values for the number of wounded nucleons is meaningful.

the grounds that the hadronization of the projectile takes place over a time-scale much longer than the time-scale of the interaction itself. In such a scenario, hadronization is a reflection of the distribution of partons in the projectile and is not sensitive to the details of the target. The increased production of pions because of target fragmentation is not of concern since these pions carry negligible energy relative to the projectile. In any case, the hadronization algorithms do not realistically handle target fragmentation even for a single nucleon target.

The second limit assumes that nuclear target effects carry over significantly to the fragmentation region though regardless of how many target nucleons participate, the final leading nucleon carries on average a significantly larger fraction of the projectile energy than the highest energy pions.

In the description of pp interactions, the initial step of the hadronization algorithm splits the available energy into a leading particle and a portion which goes through presplittings with the resulting fragments recursively generating pions. The term

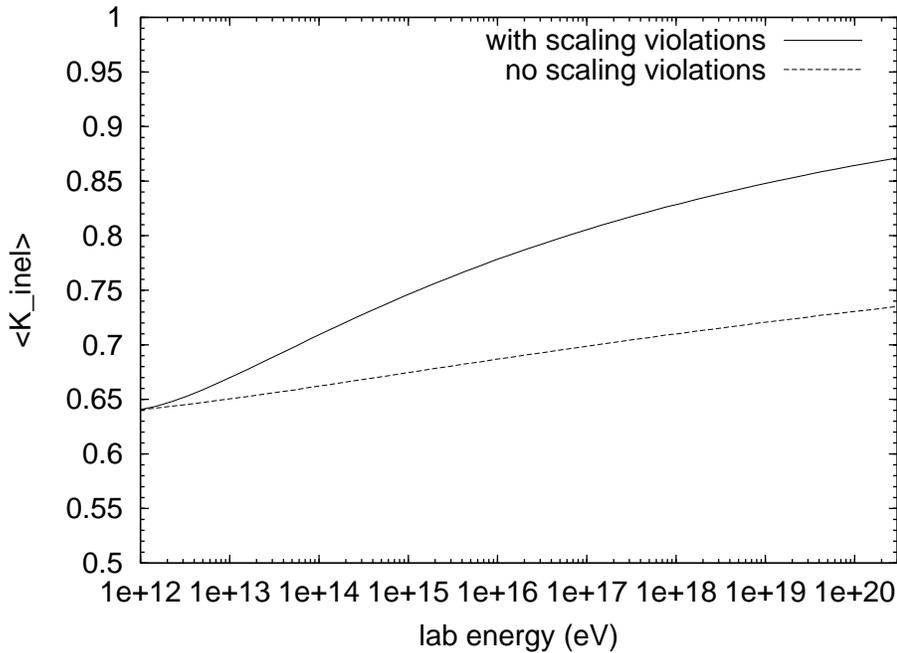


Figure 3.11: The energy dependence of the inelasticity for protons on air. The upper curve includes scaling violations in the underlying pp interactions while the lower curve does not.

presplit is adopted to describe that fraction of energy which goes through presplittings. The term *fragment* refers only to those pieces which are in the recursive phase of generating pions. A simple way of modeling the second limiting case is to recursively generate a presplit for each wounded nucleons with the exception of the last. The recursive procedure samples a distribution which takes on average a fraction of the remaining energy. For the last wounded nucleon, the remaining available energy is used to generate a leading particle and a presplit in the same manner as in pp interactions.

The average fractional energy taken by a presplit is fixed by requiring the first generated presplit to take on average the same fractional energy as the presplit generated with the last wounded nucleon. This is expressed by the relation,

$$F = (1 - F)^{N_w - 1} K, \quad (3.31)$$

where F is the average fraction of the remaining energy carried away by a presplit.

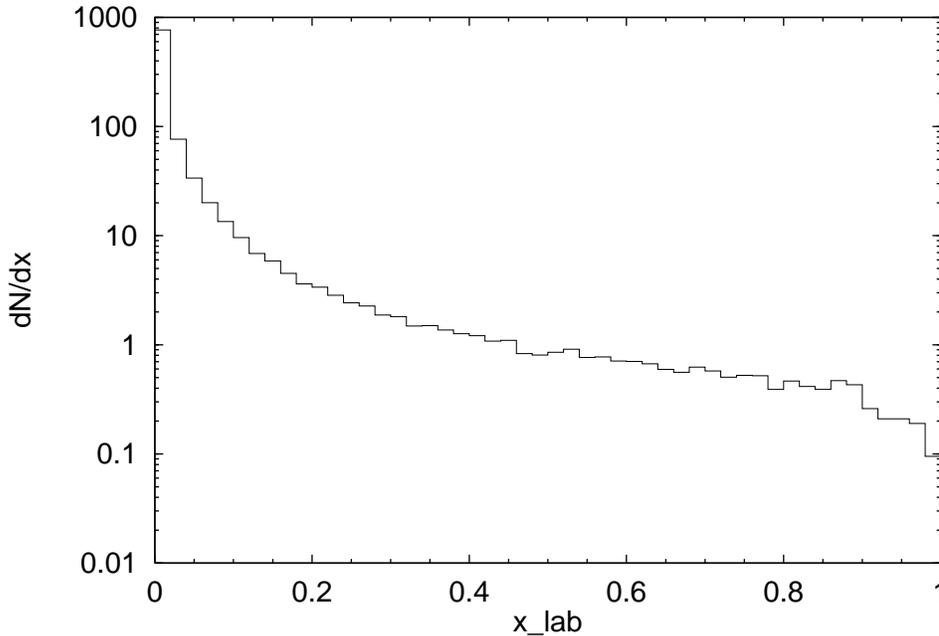


Figure 3.12: x_{lab} distribution of charged pions for charged pion projectiles on a nucleon target. The histogram is constructed from 10,000 simulated interactions.

In this way, the leading nucleon on average has the largest energy and an increasing inelasticity is accompanied by a decreasing number of energetic secondaries. All presplittings are handled with the average number which characterizes pp interactions at the projectile energy.

Equation 3.31 has been solved numerically and results have been tabulated for inelasticities in the range $[0.5, 1.0]$ at a level sufficient for interpolation. A nitrogen target is assumed. Figure 3.10 shows the results for a number of inelasticities. Figure 3.11 shows the energy dependence of the inelasticity for protons on air in the two limits where scaling violations are ignored and maximized in pp interactions.

3.7 Pion projectiles

The differences in hadronization between proton and pion projectiles are connected to the inelasticity which is higher for pions by virtue of the lower number of valence quarks, see Equation 3.17. If the same value of C in Equation 3.20 is adopted for

pion interactions as well as proton interactions, the low energy average number of presplittings is slightly larger than in proton interactions. This results in a slightly softer distribution of secondaries which is in rough agreement with fixed target data on pion-proton collisions at 250 GeV [34, 46]. Figure 3.12 shows the x_{lab} distribution of secondaries calculated in the same manner as for protons with the exception that the leading particle is also included.

Scaling violations should evolve in parallel with those in pp interactions. The rise in the average number of presplittings in pp interactions is adopted to also describe pion interactions. The constant difference in the average number of presplittings between protons and pions is that arising from the differences in the inelasticity at low energy. Such an approximation is sufficient in the context of the limiting models described above.

3.8 Nuclear projectiles

At the highest energies, primaries are unlikely to be nuclear. Interactions on the microwave background radiation break up nuclei at energies lower than those at which proton energies are degraded by the photo-production of pions (the GZK effect). The center of mass energy required needs only be sufficient to break nuclear bonding. Still, since the origin of the highest energy cosmic rays is not understood, the possibility of nuclear primaries should be considered. In general, it is best to know how all kinds of shower develop in the hopes that a comparison with experiment can give an indication about composition.

The general procedure for handling nuclear projectiles is to view the nucleus as a collection of nucleons which equally share the primary energy where each nucleon interacts inelastically with an air nucleus as described above. The initialization of an air shower simulation for a nuclear projectile entails specifying the depths of first interaction for each of the primary nucleons. A procedure often adopted for the sake of simplicity is to ignore the correlations between the depths of first interaction, taking the full shower to be a superposition of A independent showers where A is the atomic mass number of the primary. While such a procedure drastically underestimates

fluctuations in shower development, it is suitable for purposes where only average properties are under investigation. This model is referred to as the superposition model; models which treat the hadronization associated with each nucleon independently but allow for correlations in their depths of first interaction are referred to as semisuperposition models.

3.8.1 The wounded nucleon picture

An approach which handles fluctuations in shower development treats a nucleus-air interaction by grouping the nucleons emerging from an interaction into participants and spectators. Participants interact inelastically while the spectators can emerge as individual nucleons or smaller nuclei. Regardless of how the grouping of spectators is handled, the average number of participants satisfies a general relation derived from multiple scattering theory,

$$\langle N_{\text{part}} \rangle = \frac{A\sigma_{p\text{-air}}}{\sigma_{A\text{-air}}}. \quad (3.32)$$

The formalism from which this is derived is the same wounded-nucleon picture that was utilized above to treat nuclear target effects [43]. Consequently, the number of participating nucleons is sampled in the same manner which emphasizes peripheral collisions as described above.

A consequence of Equation 3.32 is that the distribution of the depths of first interaction constructed by breaking up many identical nuclei is only a function of the underlying p -air cross section. The superposition model which treats the nucleons independently is the limiting case where the number of participating nucleons is fixed at one and $\sigma_{A\text{-air}} = A\sigma_{p\text{-air}}$. This is why the superposition model is desirable in the case where interest is only in average quantities. Furthermore, since the only effect of modeling nuclear breakup in more detail is on fluctuations in development, it is suitable to adopt any rough characterization which captures general properties.

3.8.2 Nuclear cross sections

A first step in building a model which introduces correlations in the depths of first interaction is to adopt a form for nucleus-air cross sections. It is best to cast this in terms of the underlying nucleon-air cross section so that the nuclear breakup model adjusts to any change in the extrapolation of the latter to high energies. A standard low energy parameterization [47],

$$\sigma_{A_1 A_2} = \pi R_0^2 (A_1^{1/3} + A_2^{1/3} - \delta)^2, \quad (3.33)$$

with $R_0 = 1.47\text{fm}$, can be adapted to describe higher energies. This form treats the cross sections as geometric where the factor δ is included to adjust for the strength of interactions in the case of peripheral collisions. The value of δ can be determined by requiring the above formula to give the p -air cross section evaluated at the energy per nucleon of the projectile (using $A_1 = 1$ and $A_2 = 14.5$). While obviously a crude approximation, it varies with the energy dependence of the underlying p -air cross section in a physically motivated way. Regardless of how the nucleus-air cross section is parameterized, the same inclusive distribution of the depths of first interaction results.

3.8.3 The grouping of spectators

There are two limiting models of how to group spectators: either all spectators emerge as independent nucleons or all are grouped as a single nucleus. In the former case, fluctuations are underestimated, but not so much as in the superposition model. In the latter case, fluctuations are overestimated. Choosing one or the other, grouping the spectators as a single nucleus is more suitable since this is more in accord with what is expected in peripheral collisions. A better alternative might be to randomly choose one extreme or the other. The work of Engel [43] which describes the formalism of the semisuperposition model in detail suggests that a suitable choice has spectators emerging as independent nucleons 25% of the time. The differences between this mixed model and the extreme model of always grouping spectators together is small enough that a modification to include spectators emerging as multiple nuclei

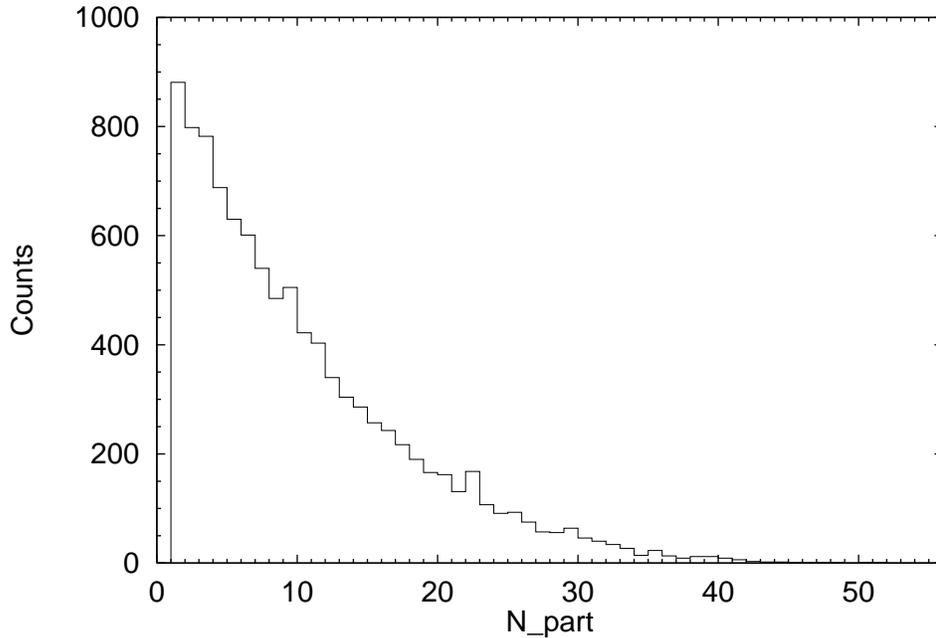


Figure 3.13: The distribution of the number of participating nucleons from the projectile for iron at 1 TeV on air. The simulation set is comprised of 10,000 interactions.

introduces an unnecessary complication in light of the crudeness of other aspects of the model.

3.8.4 Some simulation results

In order to compare this simplified version of the semisuperposition model to the more detailed results of Engel, Figures 3.13 and 3.14 show results from the simulation of iron projectiles on air at 1 TeV total energy in the lab frame. Figure 3.13 shows the distribution of the number of participating nucleons from the projectile. As expected, the distribution peaks for one participating nucleon reflecting the dominance of peripheral collisions. Figure 3.14 shows the average first depth of interaction distribution where an average is calculated over the depths of inelastic interaction for the nucleons of each iron nucleus. To the extent that fluctuations associated with the nucleon induced subshowers are negligible, this distribution is nearly the same as the depth at maximum distribution for iron-induced showers with the exception that for

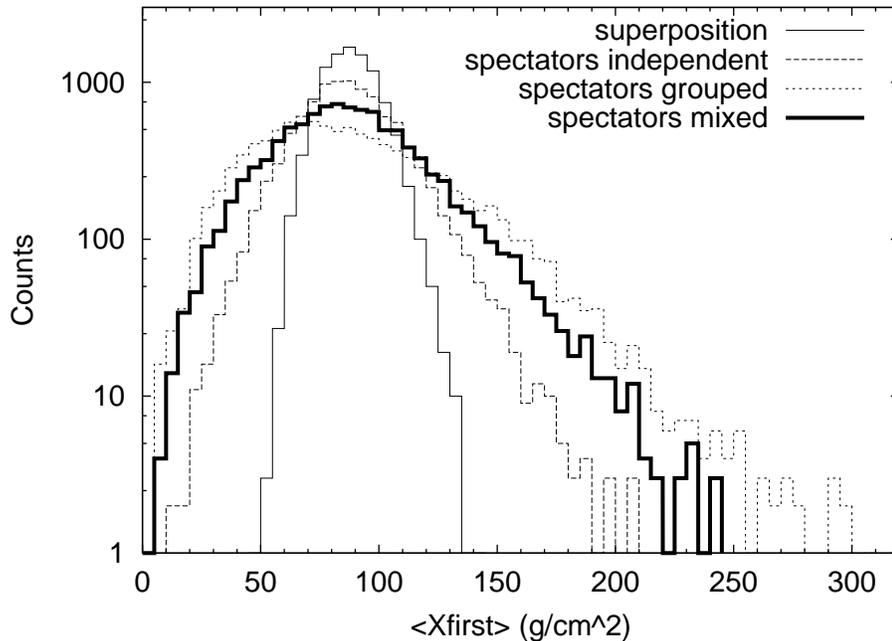


Figure 3.14: The distribution of the average depth of first interaction. Each simulation event breaks up an iron nucleus of total energy 1 TeV into nucleons each at their depth of first interaction. The simulation set is comprised of 10,000 such events.

the latter the entire distribution will be shifted to a deeper depth. If the primaries are of a higher energy, the width of the distribution is narrower since the nucleon interaction length decreases with increasing energy.

3.9 Summary

Models have been built which describe hadronic interactions for nucleon and pion projectiles on air. Also a model for breaking up nuclear primaries as they propagate through the atmosphere has been developed. Scaling violations in the hadronization phase for interactions on nucleon and air targets have been described in two limiting cases which conservatively bracket the uncertainties in extrapolation to the highest energies. This allows an assessment of model dependence and its impact on air shower development. This in turn aids in understanding the degree to which information about the primary composition or the physics associated with air shower development

can be extracted from air shower measurements. It also helps to identify what constitutes anomalous air shower development. An observation which falls clearly outside the regime of the bracketed physics described here would likely point to radically new physics or exotic types of primaries. These issues are taken up in the following chapters.

Chapter 4

Bootstrapping Models

Now that electromagnetic cascading and hadronic interactions have been described, it is time to move on to the simulation of air showers with a hadronic core. This task breaks down into two major components:

1. the development of energy dependent parameterizations which describe the various types of hadronic subshowers;
2. the simulation of full air showers with particular focus on the accurate modeling of fluctuations.

This chapter addresses primarily the first component with a careful consideration of a single baseline model of the hadronic physics: the limiting scenario where scaling violations in pp interactions and nuclear target effects have negligible impact on the forward fragmentation region. This choice is simply because this model is less computationally intensive than those with scaling violations or nuclear target effects. Also the default extrapolation of the pp inelastic cross section with $a = 0.21$, see Equation 3.9, is adopted.

4.1 The atmosphere

While interaction lengths for electromagnetic and hadronic processes are expressed in terms of the amount of matter traversed, the decay of pions takes place in time (or

equivalently distance, taking all particles to be moving at approximately the speed of light). In a medium of uniform density, matter traversed and time are the same up to a proportionality constant. In an atmosphere where the density is a function of altitude it is necessary to switch back and forth between matter traversed and time so that decisions as to whether pions interact or decay can be made. The general procedure for making such a decision on an event by event basis is to randomly sample a depth of interaction and a depth of decay, and then choose whichever process occurs first. This requires a model of the atmosphere so that the amount of matter traversed can be converted to a time and vice versa. A simple parameterization of the atmospheric density as a function of altitude, good to about 3%, is

$$\rho(h) = \begin{cases} (1.255 \times 10^{-3} \text{g/cm}^3) \exp(-h/9.192 \text{km}), & h < 10 \text{km} \\ (1.944 \times 10^{-3} \text{g/cm}^3) \exp(-h/6.452 \text{km}), & h \geq 10 \text{km} \end{cases} \quad (4.1)$$

where the column density is in units of g/cm^2 and height is in units of km [48]. For heights less than zero, it is assumed that the medium is water with a density of 1g/cm^3 . In such cases it is important to be aware that the conversions between column density and distance can easily lead to a floating point overflow. This must be guarded against as it can crash the program. The relationship between the slant depth (or the column density over the path of the shower) is related to the altitude by

$$X(h) = \sec \theta \int_h^\infty \rho(h) dh, \quad (4.2)$$

where θ is the zenith angle of the shower axis with respect to the vertical. This is valid for small zenith angles, $\sec \theta \leq 0.5$. Technically, the implication of this is that the description of hadronic subshowers has not only an energy dependence but also a dependence on zenith angle and starting depth. This is discussed in more detail below.

When calculating decay processes it is necessary to account for time dilation which depends on the energy of the unstable particle. The mean free times for decay in the rest frame and lab frame, labelled τ and τ_{lab} respectively, are related by

$$\tau_{\text{lab}} = \frac{E_{\text{lab}}}{mc^2} \tau. \quad (4.3)$$

<i>type</i>	mc^2	$c\tau$
π^\pm	$139.56995 \pm 0.00035\text{MeV}$	7.804m
π^0	$134.9764 \pm 0.0006\text{MeV}$	25.1nm

Table 4.1: Summary of pion properties needed for the calculation of decay depths in the atmosphere.

Adopting any order of magnitude density for the atmosphere and considering the mean free times of decay for charged and neutral pions, see Table 4.1, it is evident that most charged pions interact at energies above about a TeV and that neutral pion interaction becomes important at energies above about 10^{19}eV . An immediate consequence of this is that the bulk of the total shower energy is ultimately channelled into the electromagnetic cascade through the decay of neutral pions rather than in the form of muons and neutrinos stemming from the decay of charged pions. Still, while muons do not carry the dominant portion of the total shower energy, their numbers are still large. Even ground arrays which measure muon numbers only can be used to measure the cosmic ray energy spectrum [13].

4.2 Basic methodology

As discussed in the context of pure electromagnetic cascades, the general procedure adopted here for simulation of air showers is to express a total shower as a superposition of a set of subshowers each of which only makes a small contribution to the track length integral. Unlike with pure electromagnetic cascading, hadronic cascading does not scale with the energy of the primary. Even in the case where hadronization is assumed to scale approximately, there is still the energy dependence of cross sections. The assumption of scaling in the hadronization algorithm is only an approximate way of expressing that the forward fragmentation region is not significantly affected by a central region which gives high multiplicities but accounts for only a small portion of the energy. Consequently, it is quite feasible to have approximate scaling in the distribution of secondaries (the hadronization phase), but scaling violations as revealed in an energy dependence of total cross sections (the propagation

phase). Inevitably, the description of hadronic subshowers is energy dependent and thus developing appropriate parameterizations is a more involved task.

The basic methodology is the standard one used when developing parameterizations for the computation of longitudinal profiles: begin at low energy and work up to high energy making use of the lower energy parameterizations along the way [27]. Parameterizations are developed for a number of energies at a level suitable for interpolation. This bootstrap approach must simultaneously build up parameterizations for all types of hadrons in the cascade.

Because of the way the hadronization algorithms of the previous chapter have been formulated, it is useful to define a pseudo-particle which represents a collection of hadrons that can be parameterized as a single unit. The term *fragment* is used to refer to the collection of pions emerging from the iterative procedure $fragment \rightarrow fragment + pion$. From the standpoint of subshower parameterization, there is no reason not to view this collection of hadrons as a pseudo-particle. In this way hadronization of a fragment can stop as soon as the remaining fragment energy drops below a simulation threshold signifying a small contribution to the track length integral. It is tempting to extend this idea to the collection of pions referred to as a *presplit*, those objects yet to go through the presplitting phase into fragments. Unfortunately, the average number of presplittings to be conducted is not in general uniquely defined as a function of the presplit energy. Consequently all presplits must be broken down into randomly sampled fragments before parameterization can begin. This makes models which include scaling violations in pp interactions or nuclear target effects more computationally intensive to bootstrap. There are three types of parameterizations to be developed as a function of energy: nucleons, charged pions and fragments (neutral pions induce pure electromagnetic cascades and have been described using a modified Greisen formula).

4.3 The Gaisser-Hillas profile function

A flexible parameterization similar to that used for pure electromagnetic cascades is required. An appropriate choice differs from the modified Greisen form in that

it gives a steeper rise to the depth at maximum, however evolution after the depth of maximum is similar to that for pure electromagnetic cascades. A popular choice based on simulations of proton showers is the Gaisser-Hillas profile function [49],

$$N_e(X) = N_{\max} \left(\frac{X - X_0}{X_{\max} - X_0} \right)^{\frac{X_{\max} - X_0}{\lambda}} \exp \left[\frac{X_{\max} - X}{\lambda} \right] \quad (4.4)$$

where N_{\max} and X_{\max} are the size and depth and maximum, X_0 is the depth of the first inelastic interaction, and λ is a characteristic interaction length. In the fitting of simulated profiles and experimental signals, the four free parameters can not all be fitted meaningfully. In the case of experiment, usually λ is held fixed at 70g/cm^2 . This value is derived from simulations of proton showers with a scaling hadronic model. In general λ is energy dependent, but when scaling violations are present it is unclear what value for λ is most appropriate. In any case, the fit is insensitive to the choice since X_0 is a free parameter which can not be measured experimentally. This implies that a fitted X_0 with $\lambda = 70.0\text{g/cm}^2$ does not correlate strongly with the actual depth of first interaction. In the development of subshower parameterizations, the starting depth, X_0 , is normalized to zero and depths are interpreted as relative to the starting depth; λ is left as a free parameter to be fitted.

It is advantageous to use a form which gives the average profile relative to the depth of first interaction since this allows the starting depth of the subshower to be sampled randomly. Nucleons and pions are propagated before they are parameterized (fragments can not be meaningfully propagated since they represent collections of pions). This is desirable because propagation is much less computationally intensive to simulate than subshower parameterization. Also, starting depth fluctuations represent the dominant source of fluctuations in a subshower. When constructing an average profile, this means that individual subshowers should be normalized such that $X = 0$ corresponds to the depth of first interaction before adding to the total. Rather than fit N_{\max} directly, it is better to first scale the shower size down by a factor of $E_{\text{dep}}/\text{GeV}$ where E_{dep} is the average energy channelled into the electromagnetic cascade. With this convention, all showers give a normalized size at maximum, labelled below as S_{\max} , which is of order one. Finally, since the deposited energy is proportional to the track length integral, it is adopted as a fundamental parameter.

Once the deposited energy, E_{dep} , is fixed, λ is then chosen for a fixed size and depth at maximum to give the correct track length integral. The final fitting form is then,

$$N_e(X) = S_{\text{max}}(E_{\text{dep}}/\text{GeV}) \left(\frac{X}{X_{\text{rel}}} \right)^{\frac{X_{\text{rel}}}{\lambda}} \exp \left[\frac{X_{\text{rel}} - X}{\lambda} \right] \quad (4.5)$$

where the designation X_{rel} is a reminder that this is the depth at maximum relative to the starting depth of the subshower and λ is a function of $(S_{\text{max}}, X_{\text{rel}}, E_{\text{dep}})$.

4.4 Starting depths and zenith angle

As mentioned above, due to the presence of particles which can decay, the average profile for subshowers of a specific type and energy is also dependent on the starting slant depth and zenith angle. This dependence is mild since at most energies, pions of a particular type either almost always decay or else interact. Consequently, the approach taken here is to adopt characteristic values for the starting depth and the zenith angle.

For the starting depth it is best to pick a depth at which particles of the type which initiate the subshower are most likely to be found. Keeping in mind that hadronic showers evolve in a manner similar to pure electromagnetic cascades and recalling the arguments that led to Equation 2.23, an approximate value for the starting slant depth relative to the top of the atmosphere is

$$X_0 \approx X_{\text{rad}} \ln(E_0/E) \quad (4.6)$$

where E is the energy of the subshower and $E_0 \approx 3 \times 10^{20} \text{eV}$ is the primary energy of most interest. If parameterizations are developed all the way to the highest energies, a minimum starting depth can be taken to be X_{rad} .

Zenith angle is fixed at $\cos \theta = 0.5$. Roughly, beyond this angle the above description of the atmosphere needs modifications to account for the curvature of the Earth. Lesser angles are characterized by lesser solid angles, so this choice seems to be a reasonable one. Zenith angle dependence is investigated below. Even if there is any significant model dependence here, it is not so important if interest is in the

comparison of hadronic models relative to one another. Furthermore, the dependence is insignificant in the simulation of fluctuated showers since the early stages of the shower are simulated directly without use of parameterization. The final shower is affected more by the distribution of subshowers relative to one another rather than in the precise form of the individual subshowers.

4.5 Some technical details

The hadronization algorithms described in the previous chapter are simplistic in that they are unaware of particle masses: all splitting is done with a uniform deviate between zero and one and no distinction is made between energy and momentum. Rather than deal with the complications of insuring that particles are on mass shell, all particles generated with energies less than about 400 MeV are ignored. The choice for the cutoff is somewhat arbitrary; in this case it is roughly equal to a typical transverse mass, $m_T^2 = m_\pi^2 + p_x^2 + p_y^2$ with the shower axis in the z-direction. Alternatively, it is roughly the energy at which the parameterization for neutral pions reaches maximum after one radiation length. The program keeps track of all energy ignored in this manner along with the energy deposited into the electromagnetic cascade and the energy lost in the form of muons and neutrinos. This crude procedure is a valid one to the extent that the total fractional energy ignored at high energies is negligible. Note that the energy ignored due to simulation error is a parameter which is tracked at every energy for which parameterizations are developed. In this way the ignored energy is calculated as that stemming from Monte Carlo simulation at the energy under development plus that carried by parameterizations invoked during the construction of the average profile. In this way, the total fractional energy ignored at high energies is that which would be obtained if the Monte Carlo simulation were conducted without the use of parameterizations.

Bootstrapping begins at the energy minimum and extends up to 10^{21} eV over 100 energy values in uniform logarithmic intervals. With three particle types (nucleons, pion and fragments) this amounts to 300 subshower parameterizations per model. The longitudinal profile is recorded as an array of 50 uniformly spaced depths beginning

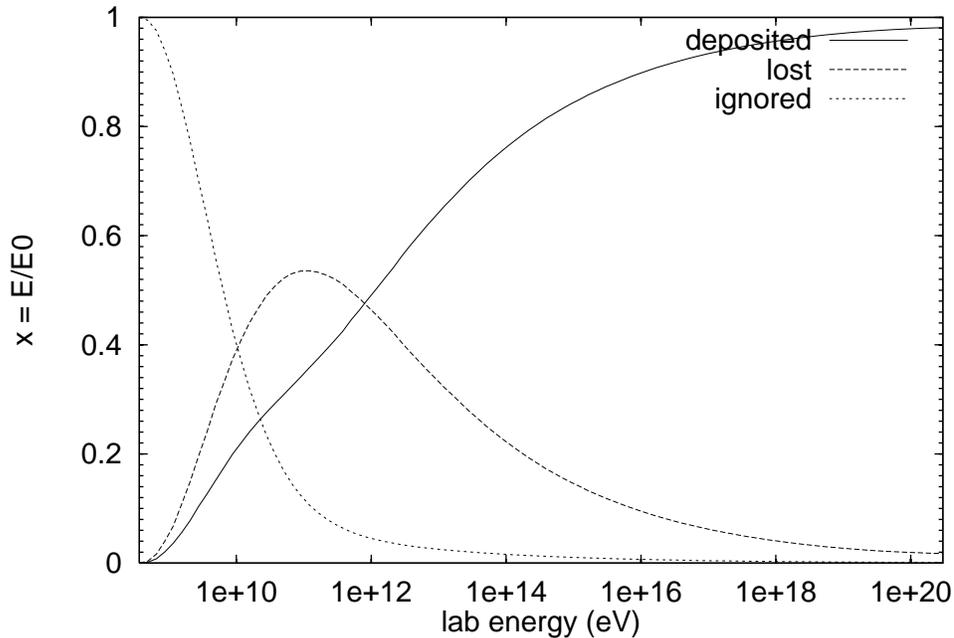


Figure 4.1: The energy dependence of the fractional energy parameters for nucleon induced subshowers in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. The figure shows the energy deposited into the electromagnetic cascade, the energy lost in muons and neutrinos and the energy ignored due to error in the simulation algorithm.

at the depth of first interaction and extending up to approximately three times the depth at maximum at which point the shower has fallen to a negligible size. So that the fitting algorithm is stable at low energies, the longitudinal profile is tracked up to at least a depth of $400\text{g}/\text{cm}^2$. An average profile is constructed over 10,000 simulated events where parameterizations are invoked whenever they are available. The quantity E_{dep} is directly determined by simulation and S_{max} and X_{max} are determined as the best pair coming from five independent random searches through the two-dimensional parameter space (a gradient descent method is used). The fitting program searches to maximize the area of overlap between the simulated and the fitted profile.

4.6 The fractional energy parameters

Figure 4.1 shows the energy dependence of the fractional energy parameters for subshowers initiated by nucleons. For all figures appearing in the remainder of this work, there are similar plots for pions and fragments which are not shown when they do not offer any different insights. The nucleon parameterizations are of more interest since they can also be viewed as representing total showers.

Since bootstrapping begins at the energy below which energy is ignored, the lowest energy is comprised of ignored energy only. The ignored fractional energy quickly drops to less than one percent at 10^{16} eV. The target energy is above 10^{20} eV; this suggests that the energy below which particles can safely be ignored can be increased roughly four orders of magnitude without compromising proton induced showers. This corresponds to energies well above a TeV!

This confirms that ignoring energy sampled at a value less than the adopted minimum has no effect of significance for the simulation of the highest energy showers. Even for iron induced showers where the reference energy is 56 times lower than the primary energy (the reference energy is the energy per nucleon), the ignored energy is negligible.

It is interesting to note that fixed target experiments which directly probe the fragmentation region are of little direct relevance since they are limited to projectiles of energies of hundreds of GeV. They are only important because of what they imply for higher energies. Thus it is easy to see why it is preferable to utilize a simple model which roughly agrees with data and extrapolates with few if any arbitrary parameters rather than a model which precisely fits data with little motivation for extrapolation to higher energies.

4.7 The depth at maximum and elongation rate

Figure 4.2 shows the energy dependence of the depth of maximum for nucleons and nuclei. For nucleons the depth at maximum is taken to be the sum of λ_{air} and the fitted value for X_{rel} . This quantity is nearly identical to the average depth at maximum

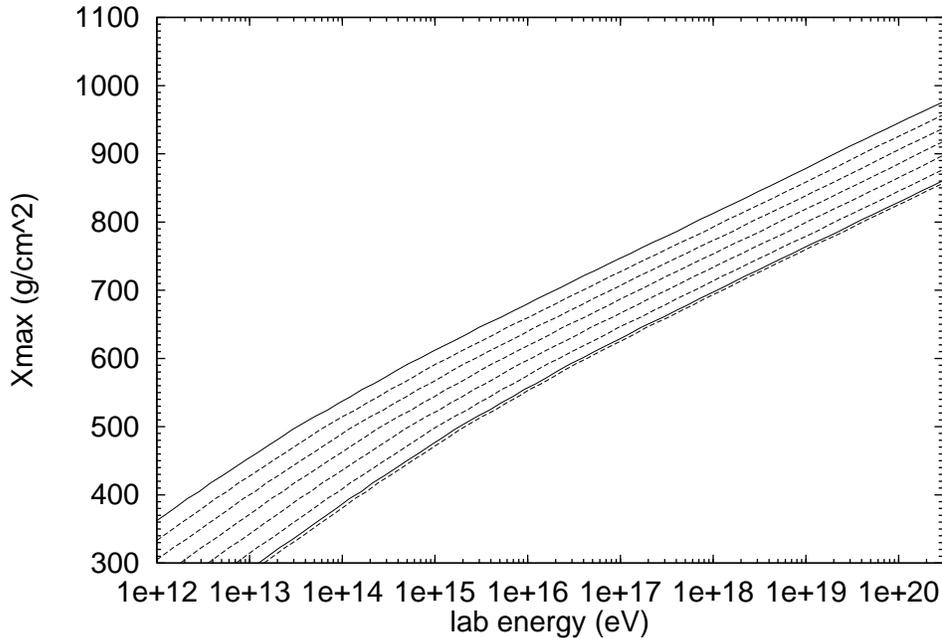


Figure 4.2: The energy dependence of the depth at maximum for nucleon and nucleus induced subshowers in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. The upper solid line corresponds to nucleons and the dotted lines correspond to nuclei with $A = 2, 4, 8, 16, 32$ and 64 . The lower solid line is for iron ($A = 56$).

over a set of fluctuated showers. The experimental determination of the average depth at maximum as a function of energy is characterized by systematic errors that have been estimated at approximately $\pm 25 \text{g/cm}^2$. This renders the distinction between the depth at maximum of the average profile and the average depth at maximum over the set of individual showers as insignificant in a study of air shower physics. The curves for nuclei are created by scaling up the energy by a factor of the mass number.

Notice that the systematic errors associated with the depth at maximum are comparable with the differences between protons and iron. This makes an analysis of the mass composition a difficult task even without the problem of model dependence. Consequently, the analysis of the mass composition by the Fly's Eye group relies on the elongation rate, defined as $dX_{\text{max}}/d \log(E)$. This slope is insensitive to a

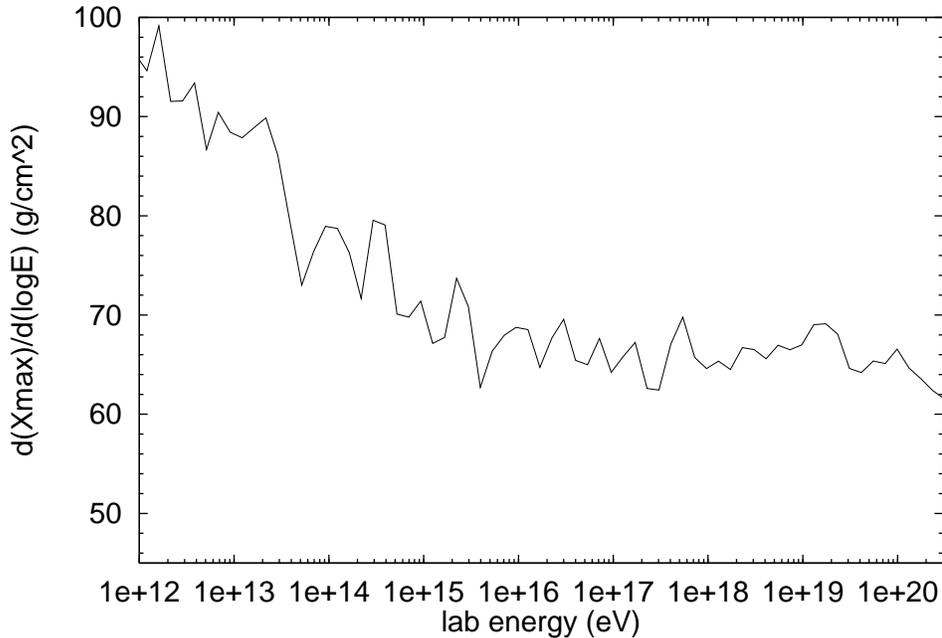


Figure 4.3: The energy dependence of the elongation rate for nucleon induced showers in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections.

systematic error over all showers.

Figure 4.3 shows the elongation rate for nucleons. In the region addressed by the Fly’s Eye experiment (roughly $10^{17} - 10^{19}$ eV), the elongation rate is approximately the same for nucleons and all types of nuclei under the assumption of a constant composition. The decrease in the elongation rate shown here is due only to the decrease in interaction lengths and in a sense is an upper bound (slightly larger elongation rates can result with cross sections that do not rise as quickly).

Regardless of the hadronic physics model, the elongation rate for a constant composition at high energy can not exceed that for pure electromagnetic cascades ($X_{\text{rad}} \ln 10 \approx 85 \text{g/cm}^2$). A scaling model of hadronic interactions with constant cross sections has the same elongation rate as for pure electromagnetic cascades since the total shower is ultimately reduced to a superposition of pure electromagnetic cascades. This is not entirely true: the energy of the shower must be such that virtually all charged pions interact and all neutral pions decay. Otherwise there is information

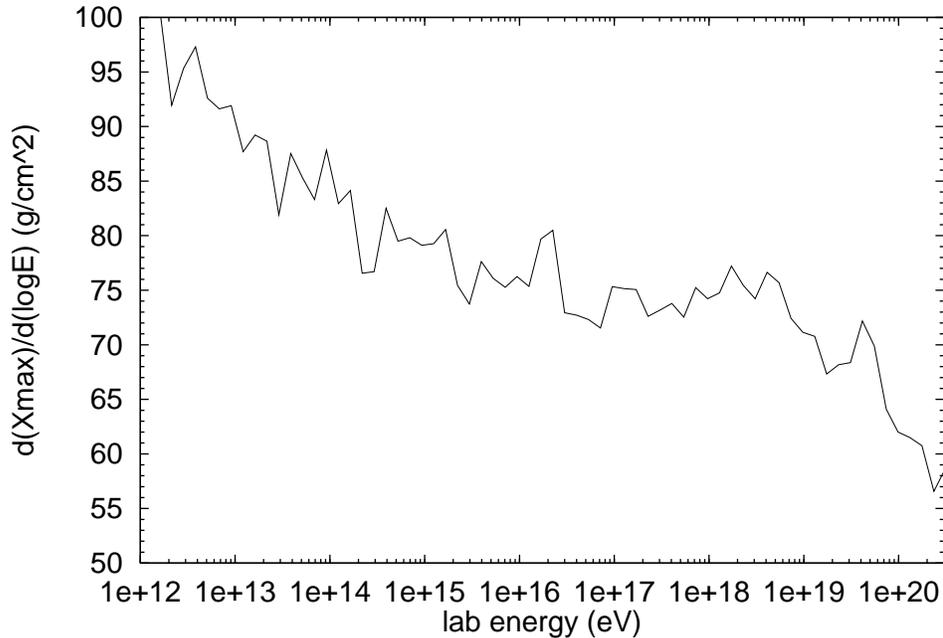


Figure 4.4: The energy dependence of the elongation rate for fragments in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections.

about the primary energy in the relative rates of decay versus interaction. This could then translate to changes in the elongation rate from that of pure scaling behavior. In fact, this is what is seen at low energies in Figure 4.3, where the effects of charged pion decay and in the errors associated with the simulation algorithm translate to an elongation rate larger than 85g/cm^2 . This is also evident at the highest energies where the elongation rate drops due to the interaction of neutral pions. This is clearly the case in Figure 4.4 which shows the elongation rate for fragments where hadronization yields higher energy neutral pions than in the case of nucleon primaries. The elongation rate at moderate energies flattens out at a higher value because there is no propagation associated with the fragment; it is still less than that for pure electromagnetic cascades due to the energy dependent interaction lengths of the secondaries derived from the fragment.

It is evident from even this baseline model that the elongation rate for a constant composition is less than that for pure electromagnetic cascades. This is the case in

any model with increasing cross sections and scaling in the forward fragmentation region. Adding scaling violations and nuclear target effects which increase multiplicities and inelasticities only decrease the elongation rate further. Examining Figure 4.2, it is evident that a composition which is dominated by heavy nuclei at energies around 10^{17} eV and becomes increasingly dominated by a lighter composition as the energy approaches 10^{19} eV implies an elongation rate which is larger than any reasonable model with a constant composition. The Fly's Eye group sees just this trend, measuring an elongation rate in this regime of 79 ± 3 g/cm². In a model independent way, there is evidence for a transition to a lighter composition.

Given that it is very difficult to come up with any mechanism which can generate the highest energy cosmic rays, it is likely that the composition consists of only a single component. In this case the elongation rate is not a useful tool for assessing composition though it could yield information about the nature of hadronic interactions as is shown in the next chapter. This also means that it is important to understand systematic errors as clearly as possible in the hopes that the measured average depth at maximum can yield information about composition. This is perhaps the most compelling reason why the Auger detector [1], for which the southern hemisphere detector is presently under construction, was chosen to be a hybrid detector. For the subset of events which are detected by the ground array and the air fluorescence detectors, there is the promise of understanding systematics at a higher level than previously possible. Redundant information can be used to cross check the detection methods.

4.8 Bootstrapping errors

This chapter closes with an investigation of some errors associated with the bootstrapping of parameterizations. Two sources of error are addressed: propagation of random error and dependence on the zenith angle of the shower axis. Both are seen to be negligible, and so the bootstrapping method as outlined here can be used with confidence.

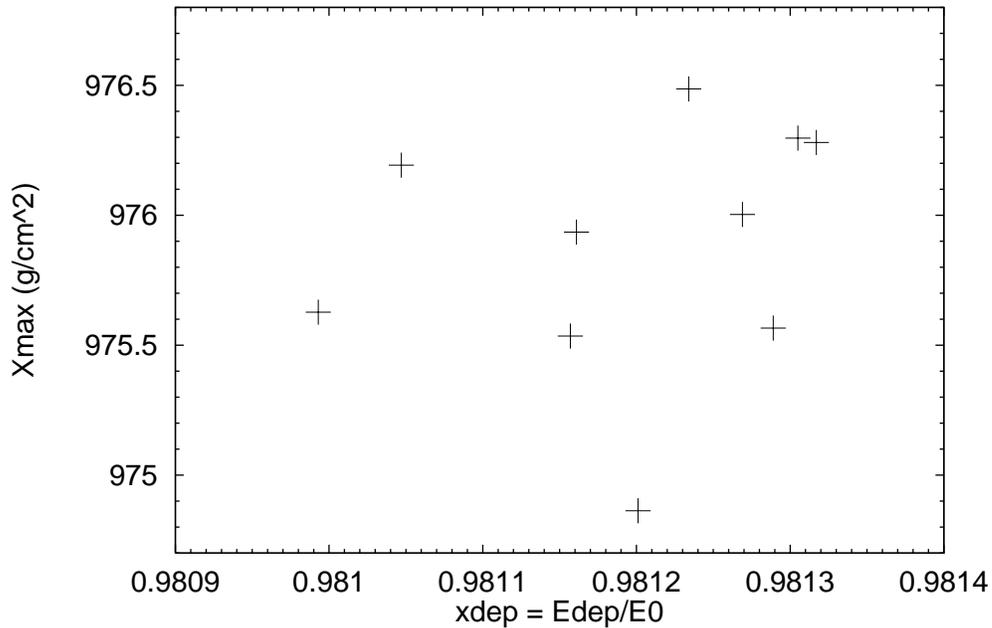


Figure 4.5: The depth at maximum and the fractional deposited energy for nucleons at 3×10^{20} eV in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections.

4.8.1 Random error

The random wandering of the elongation rates plotted above raises the question as to whether the random error associated with the bootstrapping method is significant since any error incurred at one energy during bootstrapping is propagated through all higher energies. Figure 4.5 shows the depth at maximum and the fractional deposited energy for nucleons at 3×10^{20} eV over ten independent runs (each beginning with a different seed to the random number generator). Since this model produces lower multiplicities due to the absence of scaling violations in pp interactions and nuclear target effects, the spread in values represents an upper bound for the current choice of cross sections. It is clear that the method gives the depth at maximum to within $1\text{g}/\text{cm}^2$ and the fractional deposited energy to within 0.1% ; both are negligible errors.

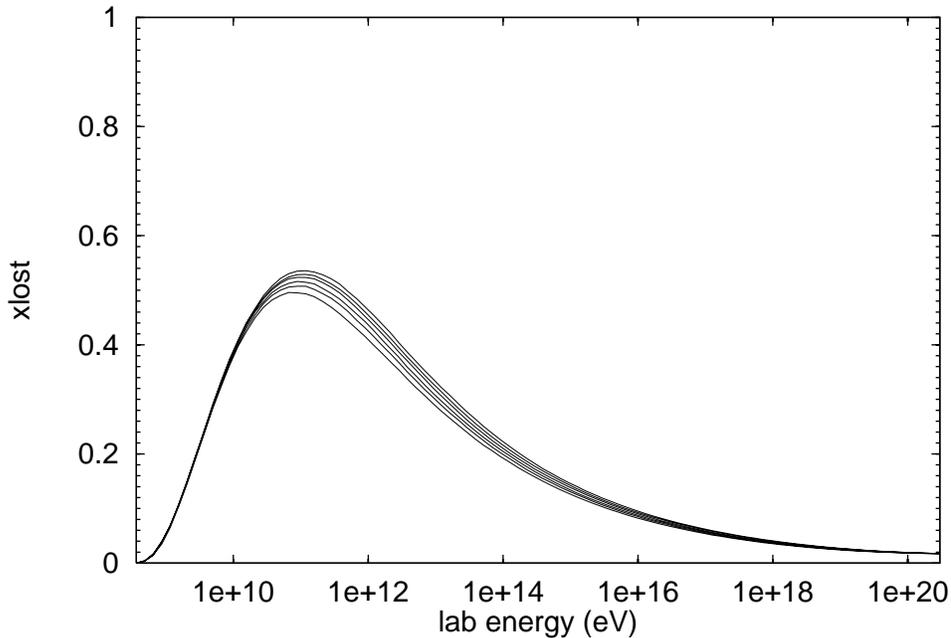


Figure 4.6: The energy dependence of the fractional energy lost in muons and neutrinos for nucleons in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. Starting from the topmost line, $\sec \theta = 0.5, 0.4, 0.3, 0.2, 0.1$ and 0 .

4.8.2 Zenith angle

The dependence on zenith angle is shown most dramatically in the energy dependence of the fractional energy lost in the form of muons and neutrinos, see Figure 4.6. Showers at large zenith angle channel a greater fraction of their energy into muons and neutrinos since they develop in the less dense upper atmosphere. However, this is not evident at the highest energies where the energies conducive to charged pion decay are many orders of magnitude away from the primary energy. The effects are inconsequential at air fluorescence energies (the threshold for the Fly’s Eye detector is around 10^{17} eV). This is seen in Figure 4.7 which shows the zenith angle dependence on the depth at maximum distribution. As one might expect, only a slight dependence is observed at energies in the range from a TeV to about 100 TeV.

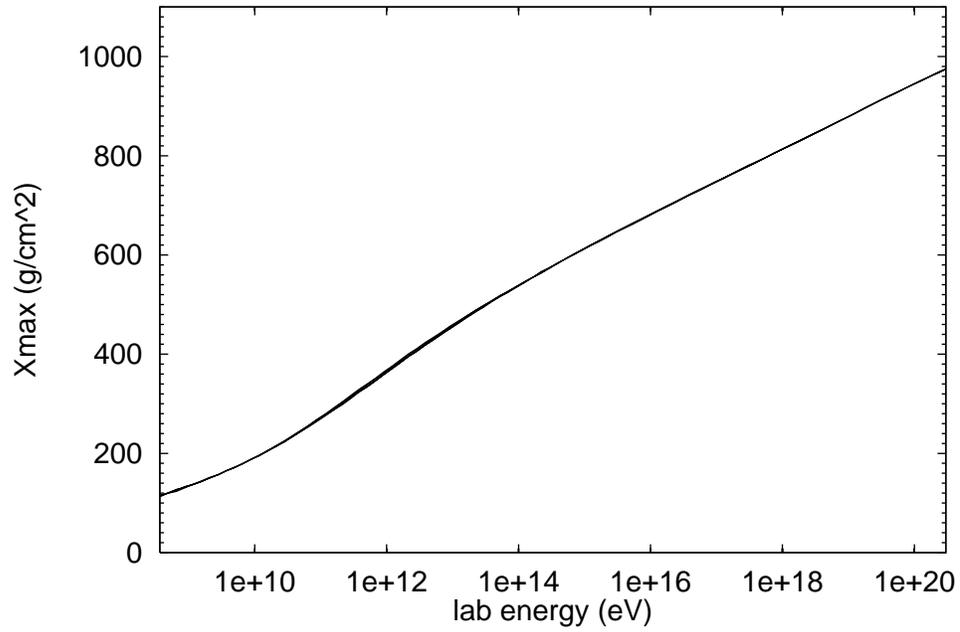


Figure 4.7: The energy dependence of the depth at maximum for nucleons in the baseline model of no scaling violations in pp interactions, no nuclear target effects and default cross sections. Depicted are six bootstrapping runs each conducted at a different zenith angle: $\sec \theta = 0.0, 0.1, 0.2, 0.3, 0.4$ and 0.5 .

Chapter 5

Model Dependence: results and implications

Since scaling violations in pp interactions and nuclear target effects have been modeled in limiting scenarios, the understanding of these hadronization models allows a clear bracketing of the properties of air showers at the highest energies. This understanding has implications for the degree to which a mass composition can be measured along with gross properties of hadronic interactions at the highest energies. This also makes clear what constitutes a signal for new physics.

Bootstrapping results for four hadronization models using default cross sections are presented and discussed. This is done by comparison with the baseline model of the previous chapter. Following this discussion is a consideration of the dependence on the extrapolation of the underlying cross sections for each of the four models. After a brief discussion of implications for measurement of a mass composition, attention is turned to the simulation of individual showers and an approximate treatment of the LPM effect. The chapter concludes with a consideration of signals for new physics in general terms and with respect to one specific possibility.

5.1 Comparing hadronization models

With limiting scenarios for hadronization on nucleon and for air targets, there are four possible hadronization models for a fixed set of cross sections. The baseline model is a scaling model of hadronization where scaling violations in pp interactions and nuclear target effects are negligible for the forward fragmentation region. The other three possible models imply showers which develop more rapidly than the baseline model. The two quantities which are explored here are the depth at maximum and the elongation rate. Figures 5.6, 5.7 and 5.8 show the results of bootstrapping for each of the three models in comparison with the baseline model.

The turning on of nuclear target effects makes an immediate difference at all energies in the depth at maximum and causes a mild drop in the elongation rate. In contrast, the turning on of scaling violation in pp interactions leaves the lowest energies unaffected but causes a more drastic drop in the elongation rate. At the highest energies, the impacts of nuclear target effects and scaling violations are comparable. When both effects are turned on, as one might expect, there is an immediate drop in the depth at maximum due to the turn on of nuclear target effects and a drastic drop in the elongation rate due to scaling violations in pp interactions. When extrapolated to the highest energies, it is difficult to distinguish between nucleons with all effects turned on and iron in the scaling model. In all cases, the elongation rate for nucleons is significantly smaller than that measured by the Fly's Eye group ($79 \pm 3\text{g/cm}^2$) indicating a transition from a composition dominated by heavy nuclei at around 10^{17}eV to a significantly lighter composition at 10^{19}eV .

5.2 Comparing extrapolations of cross sections

Model dependence associated with cross sections is treated by investigating different extrapolations of the underlying pp inelastic cross section while maintaining the relationship of this cross section to all other relevant cross sections as outlined in Section 3.1. It is perhaps important to state that the debate concerning the rather uncertain connection between the inelastic cross sections for hadrons on air to the total

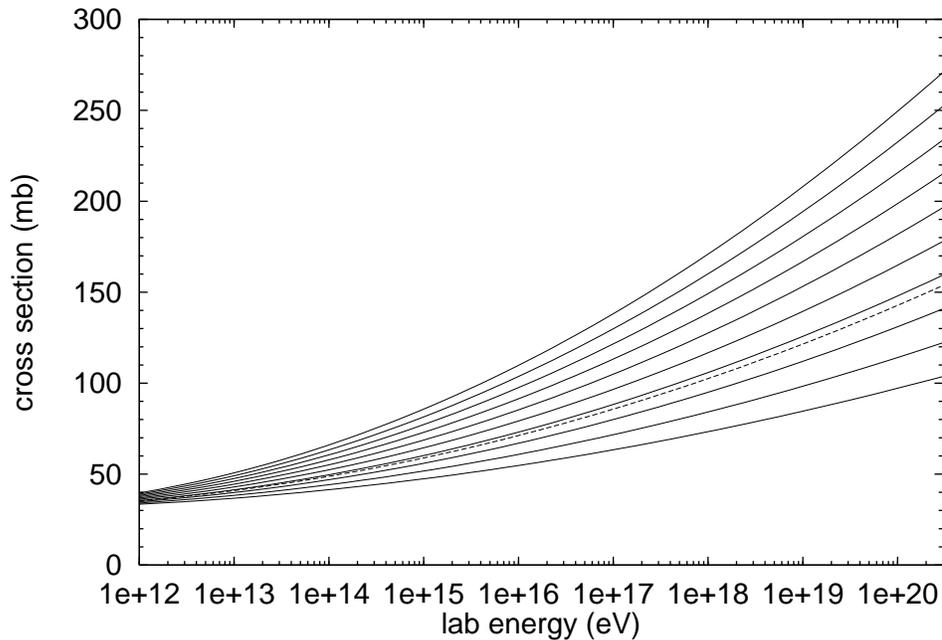


Figure 5.1: The energy dependence of the inelastic pp cross section over the set of models investigated. The dotted line shows the default.

cross section for hadrons on a nucleon target is of little concern here. The connection between the inelastic cross sections on air and on nucleon targets is only mildly model dependent [30]. The first order of business is to determine an appropriate range over which to test different extrapolations.

5.2.1 Constraining the p -air inelastic cross section

The inelastic p -air cross section is constrained roughly from the Akeno ground array experiment and from the Fly’s Eye air fluorescence experiment. Though the techniques of these experiments are quite different, both rely on an investigation of the attenuation length associated with the most deeply penetrating showers. It is assumed that for a fixed energy the most deeply penetrating showers are mostly induced by protons. The attenuation length characterizes the exponential tail of the depth at maximum distribution. It is proportional to the interaction length for protons on air though the proportionality constant is dependent on assumptions about hadronization

and on the level of contamination by helium and heavier nuclei. Consequently an analysis of the tail of the X_{\max} distribution for Fly’s Eye stereo data only weakly constrains the inelastic p-air cross section to lie somewhere in the range $530 \pm 100\text{mb}$ at $\sqrt{s} \approx 30\text{TeV}$ [23]. Using Equation 3.5, this translates to a range of $70 - 160\text{mb}$ for the underlying pp inelastic cross section. This range is bootstrapped in 10mb intervals for a total of 10 extrapolations for each hadronization model. Figures 5.1 and 5.9 show the energy dependence of cross sections and interaction lengths for nucleon projectiles.

5.2.2 Bootstrapping results

Figures 5.10 and 5.11 show the results for bootstrapping the set of cross section extrapolations for each of the hadronization models. Note that the difference in the interaction length between the two extremes at the highest energies is about 20g/cm^2 whereas the spread in the values of X_{\max} is about 100g/cm^2 . While the spread is larger by a factor of five, the relative spacing between successive models is roughly the same. This suggests that the development of the longitudinal profile is sensitive to the development features of the hadronic core over at least five generations of interactions.

The uncertainties associated with hadronization and with cross sections are comparable to each other and also to the relative difference between protons and nuclei. With regard to the Fly’s Eye measurement of the mass composition, given the large systematic errors associated with the experiment, it is not possible to rule out any of these models based on the depth at maximum at low energies for a heavy nucleus dominated composition. The scaling model with the lowest cross sections represents the lower limit for consistency with a composition of pure iron at around 10^{17}eV . The overall picture is that experimental results can be interpreted consistently with a large range of models and compositions [27, 44]. This explains the reliance on the elongation rate to make the mild but model independent assertion of evidence for a transition from a heavy to light composition in the range $10^{17} - 10^{19}\text{eV}$

5.3 Nuclear target effects revisited

The bootstrapping results indicate that the manner of modeling nuclear target effects have significant effects on shower development. Since these effects are evident even at low energies, the question is raised as to whether either of the limiting scenarios can be adopted with confidence. Evidence from fixed target experiments addressing p -nucleus collisions at beam energies of 100 GeV (for C, Al, Cu, Ag and Pb) and 120 GeV (for Be, Cu, Ag, W and U) suggest that the limiting scenario where nuclear target effects are maximized gives a reliable description.

The analysis of Fricter, Gaisser and Stanev [50] focuses on the inelasticity characterizing results in terms of the partial inelasticity coefficient defined by the relation

$$\langle E \rangle_\nu = (1 - I_\nu) \times \langle E \rangle_{\nu-1} . \quad (5.1)$$

This relates the average energy of the leading nucleon to the number of wounded nucleons from the target. It is assumed that I_1 is fixed by pp interactions, leaving $I_{\nu>1}$ as free parameters to be fitted.

The limitations of the data allow reasonable fits to be obtained if all $I_{\nu>1}$ are taken to be the same. Two limiting models of treating the isospin of the leading nucleon in the context of the wounded nucleon picture were analyzed: the first assumes that the isospin of the nucleon is determined in the interaction with only one wounded nucleon, the other considers a naive treatment where each wounded nucleon can change the isospin of the leading nucleon with equal probability. The data are fitted with values of $I_{\nu>1} = 0.18$ and 0.14 respectively for the two models, though 0.18 gives a better overall fit.

While the modeling of nuclear target effects developed in the present paper is different than the basic picture of the above analysis, Figure 3.10 indicates that the trends are similar. Characteristic values for the fractional energy carried away by presplits, see Equation 3.31, lie within the range implied by the analysis. For purposes of illustration, the analysis of Fricter, Gaisser and Stanev calculates the corresponding inelasticities for p -air interactions at low energy assuming the canonical value of 0.5 for the underlying pp inelasticity even though their fit to pp data implies a

larger value. This gives $K_{\text{inel}} = 0.63$ and 0.56 for the models corresponding to $I_{\nu>1} = 0.18$ and 0.14 respectively. The inelasticity in p -air interactions can be calculated using the model of the present work also assuming the canonical value of 0.5 for the pp inelasticity, and it is found to be 0.63 , the same as with the model which gives the better overall fit to the p -nucleus data. This suggests the model developed in this work is a reasonable one. Since the spread between the limiting scenarios is comparable to that between the spread in the limiting models of scaling violations in pp interactions, adopting the model of nuclear target effects while retaining the uncertainty associated with pp interactions should reasonably bracket the hadronization uncertainties.

Figure 5.12 shows the resulting limiting models for the depth at maximum and the elongation rate. The solid line lower bound turns on scaling violations in pp interactions and adopts the largest rise in cross sections consistent with air shower data. The middle line also turns on scaling violations in pp interactions but adopts the default cross sections. The default cross sections represent the best $(\ln s)^2$ fit to the data and also approximately represents the upper bound over the range of extrapolations which have been calculated using theoretical models. The physics which approximately gives this rise is based on an underlying theory of parton-parton scattering [36]. The upper solid line turns off scaling violations in pp interactions and uses the smallest cross sections consistent with air shower data. This choice for the cross sections is close to the lower bound over models currently in use [51].

5.4 Searching for new physics

Showers which penetrate a significant amount deeper than is implied by Figure 5.12 constitute a clear signal of new physics and/or exotic sources. For example, Domokos and Kovesi-Domokos [17] discuss a possible model of grand unification in which neutrinos acquire a strong interaction near the highest air shower energies. This unification could occur possibly around $\sqrt{s} \approx 30\text{TeV}$ or even lower and be consistent with precision accelerator measurements. In such a scenario, the neutrino inelastic cross section rises exponentially as a function of energy and then flattens out and presumably rises in a similar fashion to typical hadronic cross sections. The turn off

of the exponential rise is needed so that unitarity bounds are respected.

While the cross section is probably of the order of typical hadronic cross sections, since the neutrino is not a composite particle like hadrons it is difficult to say anything more than this. Given the uncertainties, a simple model that characterizes this exponential turn on of new physics is to treat the neutrino cross section using a step function leaving the turn on energy and the cross section as free parameters. Such a scenario is attractive since it avoids the problems associated with propagation through the cosmic microwave background radiation: because the neutrino is massless, the center of mass energy which allows for pion production is higher than that for protons. The hadronization in a ν -air interaction probably looks something like that in hadron-hadron interactions. Since the neutrino is not composite, it is likely that nuclear target effects are negligible (these effects are typically associated with the composite structure of the projectile). A scenario where the neutrino cross section is somewhat less than the nucleon cross section could easily produce longitudinal profiles which penetrate significantly deeper.

5.4.1 Muons in proton showers

It is impossible to say anything definite about hadronization at GUT energies. It is likely that all types of particles including muons are produced in comparable numbers. This suggests there may be an observable increase in the number of high energy muons due to their prompt production at high energies.

Since interest is only in general trends, the bootstrapping algorithms have been modified to track muon numbers above a number of thresholds starting at a GeV and going up to a PeV by factors of ten. While the treatment here is not suited to the needs of ground arrays, it can still be useful in the relative comparison of hadronic models and in the search for possible unique signatures of anomalous showers. Ground arrays in a rough way can discriminate between thresholds at around a GeV and about 10 GeV by exploiting timing information related to the curvature of the shower front. Deep underground, underwater and under-ice detectors can use the medium above as a filter pushing the muon threshold up to as high as a TeV. While the highest

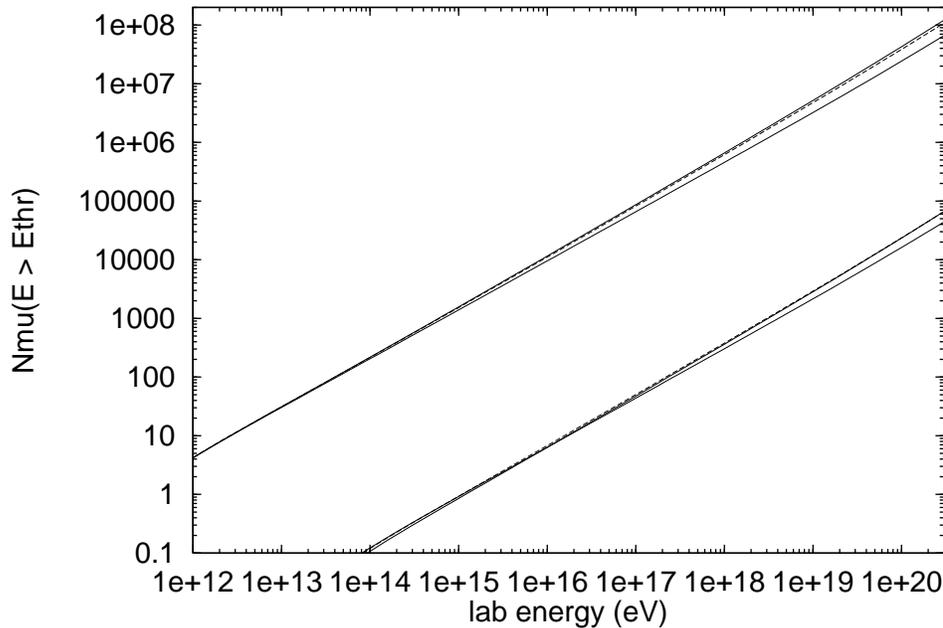


Figure 5.2: The energy dependence of the number of muons for thresholds of 10 GeV and 1 TeV evaluated for the bracketing models. Notice that the muon number is not strongly sensitive to the choice of cross sections

thresholds tracked are of no relevance experimentally, they are not a bottleneck on computation time and they can serve as useful checks for the program and for the modeled physics. Figure 5.2 shows the muon numbers for thresholds of 10 GeV and 1 TeV over the bracketing models from Figure 5.12. The result for nuclei can be obtained by scaling the energy down by a factor of the mass number and then multiplying the muon numbers at that energy by the mass number.

5.4.2 Prompt muons in anomalous showers

An order of magnitude estimate of the number of prompt muons can be calculated by counting the total number of particles produced in interactions which take place above the GUT energy for a given primary energy. The hadronization algorithm can be any algorithm of hadronic interactions, ultimately the hadronization of all particles above the GUT energy reduces to a set of fragments all with energies less than the

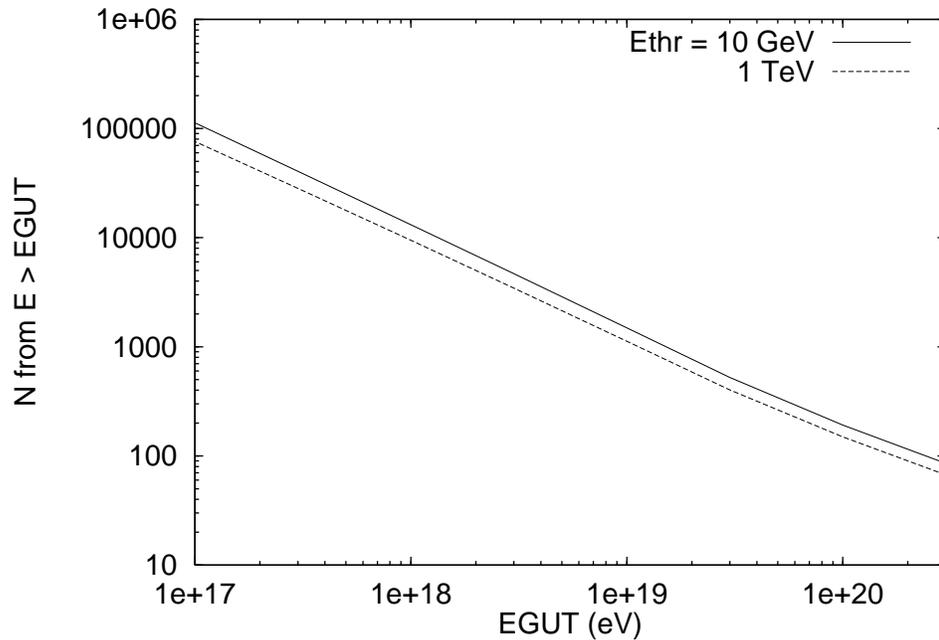


Figure 5.3: The total number of particles above a specified threshold generated in anomalous interactions above the unification energy as a function of the unification energy.

GUT energy. The resulting multiplicities are comparable regardless of how they are generated. The number of prompt muons is some fraction of this total number. The total number of muons is the sum of prompt muons plus a number of the order of but less than that for nucleon induced showers. The threshold needed for a clear signal is that for which the number of particles generated in above GUT energy interactions is comparable to the number of muons in nucleon induced showers. Prompt muons are closer to the shower axis and so the number need not be significantly larger only comparable so that in the region close to the shower axis the increased density is potentially discernible. An assessment of this requires a simulation of the lateral spread in the shower. The treatment here serves as a preliminary investigation to determine if more detailed simulations are warranted.

Figure 5.3 shows the average number of particles generated in GUT scale interactions for primary neutrinos of energy 3×10^{20} eV as a function of the approximate turn-on GUT energy. The original Hillas splitting algorithm was used to conduct all

interactions for particles with energies above the GUT energy. The total multiplicity is that resulting from the sum over all these interactions. Muon thresholds of 10 GeV and 1 TeV were simulated. These results indicate that there is no experimental signal at 10 GeV and possibly a marginal signal at 1 TeV for a low GUT threshold and good resolution of the shower near its axis where a jump in the density of muons might be discernible. It seems that the longitudinal profile offers the best chance of observing an anomalous signal.

5.5 The LPM effect

So far the discussion has focused on bootstrapping results developed under the assumption of the standard Bethe-Heitler description of electromagnetic processes. However for media of high density or for particles of very high energy, the cross sections for bremsstrahlung and pair creation are reduced because of multiple scattering (the Landau-Pomeranchuk-Migdal or LPM effect) [52, 53]. For photons and electrons in air with energies greater than about 10^{17} eV, the suppression is appreciable. For hadronic showers with total energy less than about 10^{20} eV, this effect is not significant since the highest energy electrons and photons in the cascade typically have energies much less than the primary energy. For heavy nuclei, characteristic energies for electrons and photons are even lower since the reference energy is the energy per nucleon. Even for protons at the highest energy the effect is not as significant as one might expect since neutral pions above about 10^{19} eV often interact inelastically rather than decay.

In keeping with the approximations made thus far, a simple model is considered which varies in a reasonable way with energy, atmospheric density and particle type. The features which have been modeled are taken from the discussion by Klein [54] which focuses on implications for air showers. The most important consequence of the LPM effect for air showers is an increase of the radiation length associated with electromagnetic processes. For an isothermal model of the atmosphere

$$X_{\text{rad}} \approx (37.1\text{g/cm}^2) \sqrt{\frac{E}{E_{\text{LPM}}}} \quad (5.2)$$

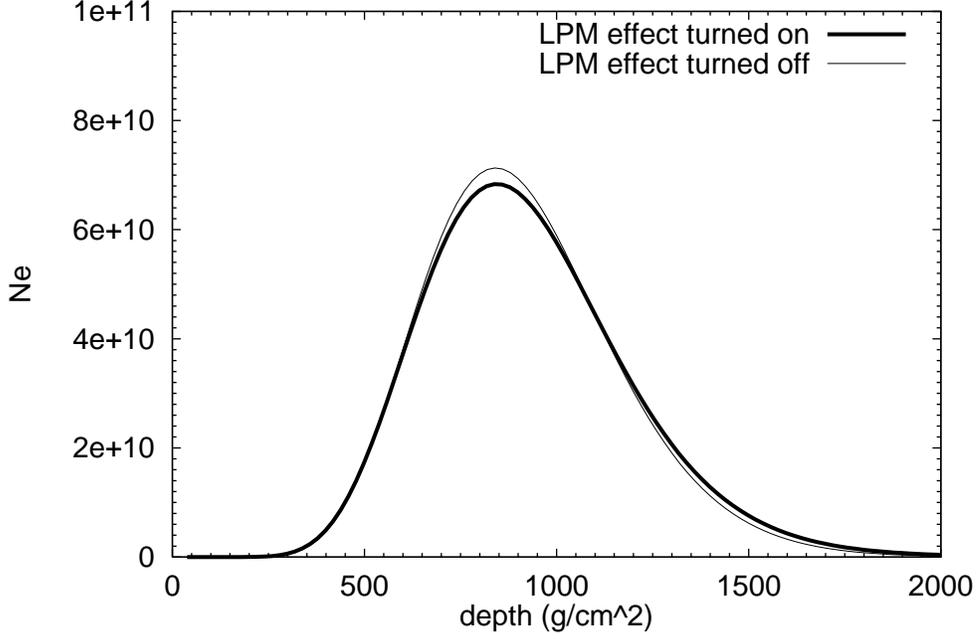


Figure 5.4: The average profile over 1000 showers for protons at 10^{20} eV. The starting depth and development phase are averaged separately so the resulting shape is more characteristic of individual showers. The hadronization model uses default cross sections with scaling violations in pp interactions and nuclear target effects turned on.

where

$$E_{\text{LPM}} \approx (117\text{PeV}) \frac{X_v}{1030\text{g/cm}^2}, \quad (5.3)$$

and X_v is the vertical depth in the atmosphere. For particle energies less than the E_{LPM} suppression is negligible and the standard value of 37.1g/cm^2 for the radiation length in air can be used. Equation 5.3 is applicable when the suppression effects are small. In the limit of large suppression, there appears an additional term in the cross sections due to a non-Gaussian tail associated with large angle scatters resulting in the form,

$$X_{\text{rad}} \approx (37.1\text{g/cm}^2) \frac{\sqrt{E/E_{\text{LPM}}}}{\sqrt{\log(E/E_{\text{LPM}})}}. \quad (5.4)$$

This form is adopted when $\log(E/E_{\text{LPM}}) > 1$ so that the transition between the two forms is continuous.

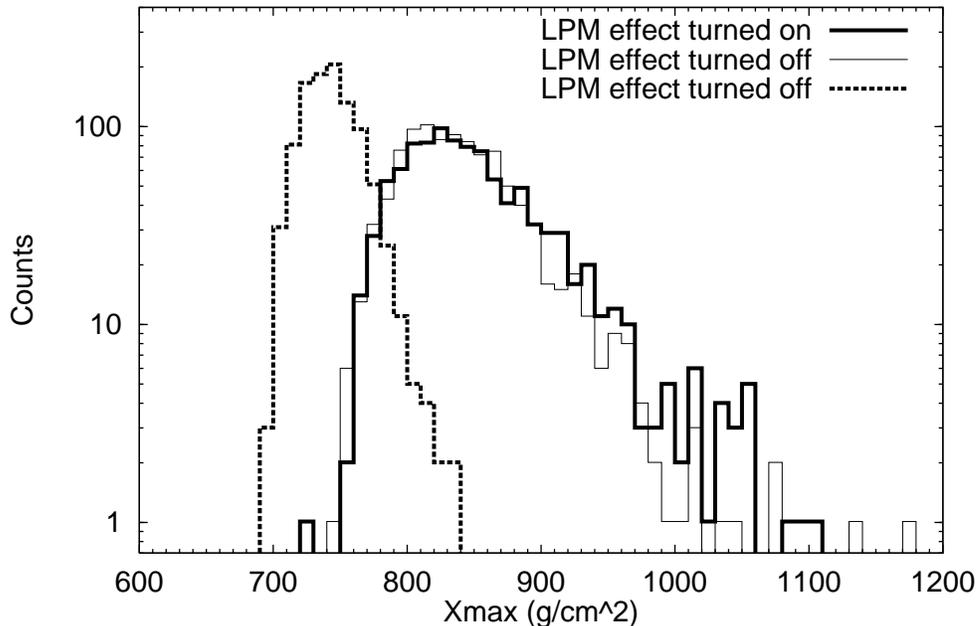


Figure 5.5: The depth at maximum distribution for protons and iron at 10^{20} eV over 1000 showers in each set. Iron is not sensitive to the LPM effect so only one histogram is shown in this case. The hadronization model uses default cross sections with scaling violations in pp interactions and nuclear target effects turned on.

In the context of a Monte Carlo algorithm even the implementation of a varying radiation length is difficult to handle since as a particle propagates, the radiation length increases. Evaluating the radiation length at the depth of the particle's creation underestimates the influence of the LPM effect on shower development. Instead, a characteristic radiation length is chosen which satisfies the relationship,

$$\lambda = X_{\text{rad}}(E, X_v + \alpha \lambda \cos \theta) \times f \quad (5.5)$$

where f is the ratio of the mean free path to the radiation length in the absence of the LPM effect (about 1.29 for photons and approximately 0.31 for electrons in the effective splitting model). In other words, the radiation length adopted is that evaluated at the depth after which the particle has propagated approximately α mean free paths, where α is a free parameter of the model (taken to be 1 below). Equation 5.5 can be solved numerically using a simple iterative procedure. A slightly longer characteristic radiation length is obtained for photons than for electrons since they stand

a higher probability of penetrating deeper into the atmosphere where suppression is larger. Similarly, the characteristic radiation length grows as the zenith angle approaches the vertical. The iterative procedure is conducted until the characteristic radiation length is determined to within $0.1\text{g}/\text{cm}^2$.

In cases where propagation goes deeper than sea level, the characteristic radiation length at sea level is assumed as opposed to water. Consequently, a lower bound on the LPM energy can be set, that for air at sea level (117 PeV). When the LPM effect is turned on in the shower code, this energy is used as the largest possible threshold even if a higher threshold is requested. This is because parameterizations are developed under the assumption of Bethe-Heitler physics and so they can not be invoked at energies where the LPM effect is potentially significant.

In relation to proton initiated showers, the LPM effect has been implemented in detail by Kalmykov, Ostapchenko and Pavlov [55]. For vertical proton showers at 10^{20}eV they find a reduction in the size at maximum by about 5%. While this is small with respect to an individual shower, this systematic effect can become important in the analysis of a set of events. Figure 5.4 shows the average profile generated with the above simplified model. Similarly, the simplified model reduces the size at maximum by about 5%. Figure 5.5 shows the depth at maximum distribution for the 1000 showers in each set used to construct the average profile. Also shown is a distribution for iron primaries at the same total energy (the same distribution for iron primaries at this energy results whether the LPM effect is turned on or off). While the LPM effect results in a wider distribution owing to the deeply penetrating tail, the effect is not so drastic relative to the differences between protons and iron.

5.6 Fluctuations and primary type

There are a few other features worth noting with respect to Figure 5.5. Comparing the distribution for iron with that obtained by averaging the depths of first interaction of the primary nucleons, the analog of Figure 3.14 conducted at 10^{20}eV , it is evident that the width and shape of the distribution is insensitive to the hadronization model for nucleons and pions. In other words, fluctuations in shower development are due to

how the nucleus breaks up in the atmosphere but not to the subsequent hadronization.

Also, allowing for model dependence it is evident that distinguishing between even the two extremes of iron and protons is impossible on an event by event basis except at the extremes of uncharacteristically shallow or deep development. Using the elongation rate to shift the distributions up to 3×10^{20} eV, it is clear that the highest energy event observed so far by the Fly's Eye group is a completely normal looking event (apart from the enormous energy). It is consistent with either a nucleus or a proton. Over sets of events, the fluctuations in shower development can be an important tool in analyzing the composition, particularly in the case of a single component composition since in such a case the elongation rate is not helpful.

5.7 Summary

The features of shower development at the highest energies have been constrained under assumptions about composition (protons or nuclei) and the extrapolation of the hadronic physics (smooth evolution consistent with general expectations connected to the structure of hadrons). Electromagnetic cascading was implemented in a manner which yields consistency between Monte Carlo simulation and parameterization. Also included was a treatment of nuclear projectile breakup which includes correlations in the depths of first interaction of the primary nucleons and a rough model of the LPM effect. The result is an air shower generator which reliably models fluctuations in shower development and can be easily adapted to describe various extrapolations of the hadronic physics either through the use of one of the many models investigated here or in the bootstrapping of new models.

In conjunction with a detector Monte Carlo, this generator is a powerful tool which aids in the development of an optimal detector design and helps uncover the potential for the next generation air fluorescence detectors to understand the composition and physics associated with the cosmic rays of the highest energy. This generator is currently being used by the OWL group. OWL is a proposed detector which will attempt to obtain a large statistics observation of the highest energy cosmic rays through a pair of orbiting air fluorescence detectors looking down into the atmosphere

from a height of about 500 km.

This generator can also be useful to theoreticians who are searching for models which can explain the events observed so far and predict features which may be revealed in future observations. A preliminary investigation of one such theory with respect to its implications for shower development was investigated above. The large uncertainties associated with hadronic physics at the highest energies require very long length scales to be associated with new physics or exotic sources to clearly discriminate from the traditional compositions of protons and nuclei.

All in all, the uncertainties associated with the physics investigated above in conjunction with the systematic and random uncertainties associated with experiment will make the extraction of a mass composition or the features of hadronic interactions at the highest energies a formidable task. Most likely, analysis will be done in a way which constrains the possible pairings of composition and interaction physics. Fortunately, the measurement of the energy spectrum is a much less tenuous undertaking and the shape of the measured energy spectrum will likely contain information about the kinds of mechanisms responsible for the existence of these highest energy cosmic rays observed at the Earth. Once a high statistics observation is obtained, a global analysis which takes into account the measured energy spectrum, constraints from a study of composition and interaction features, and also theoretical models which explain acceleration to the highest energies and propagation to the Earth will hopefully yield a consistent picture. Even at the modest level of observations obtained so far, such a consistent picture is elusive.

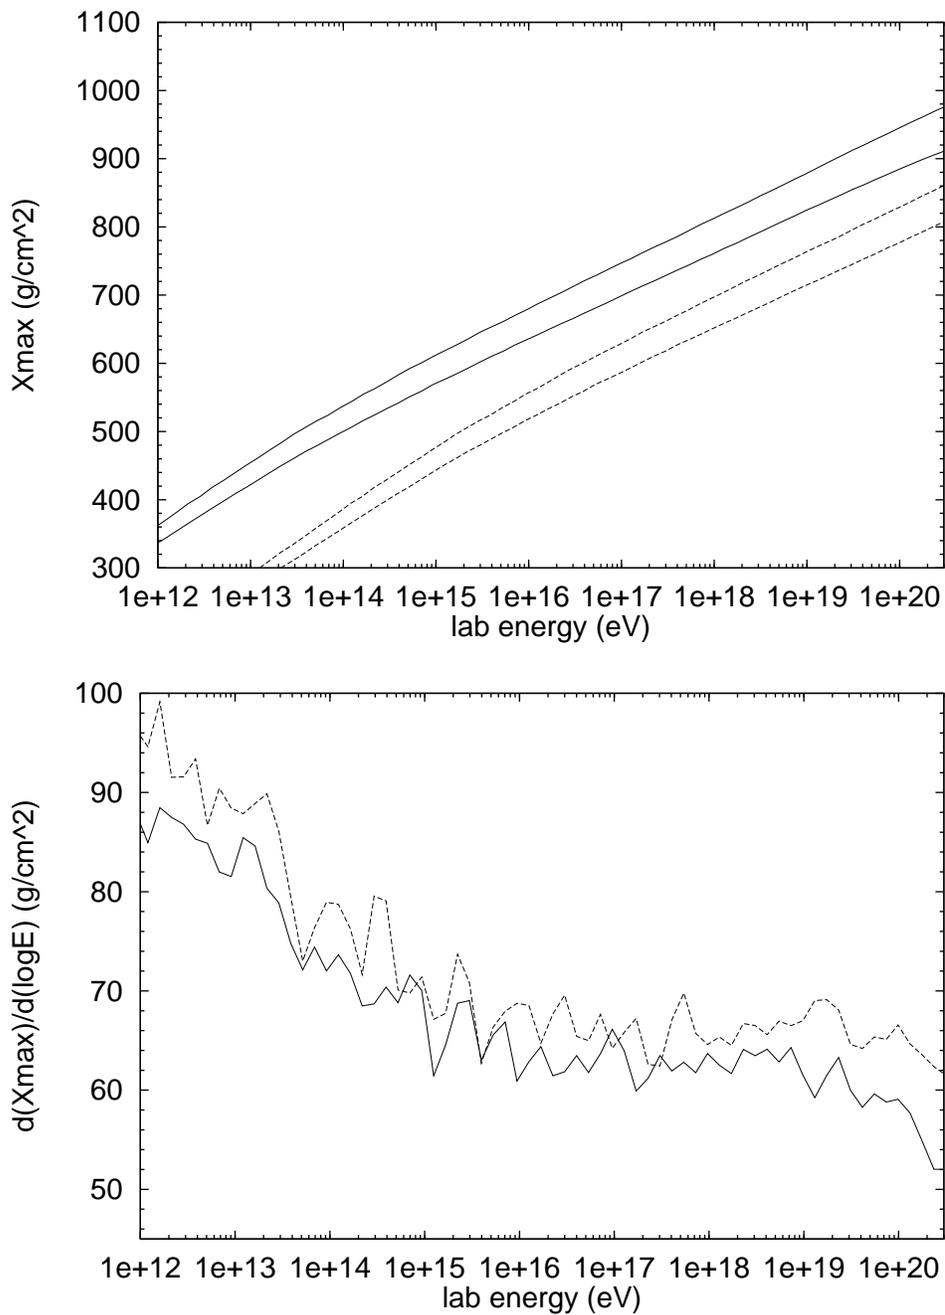


Figure 5.6: Using default cross sections, the energy dependence of the depth at maximum and the elongation rate for nucleons with only nuclear target effects turned on. The dashed lines in the depth of maximum plot are for iron. The results are compared with the baseline model, the upper line of each pair.

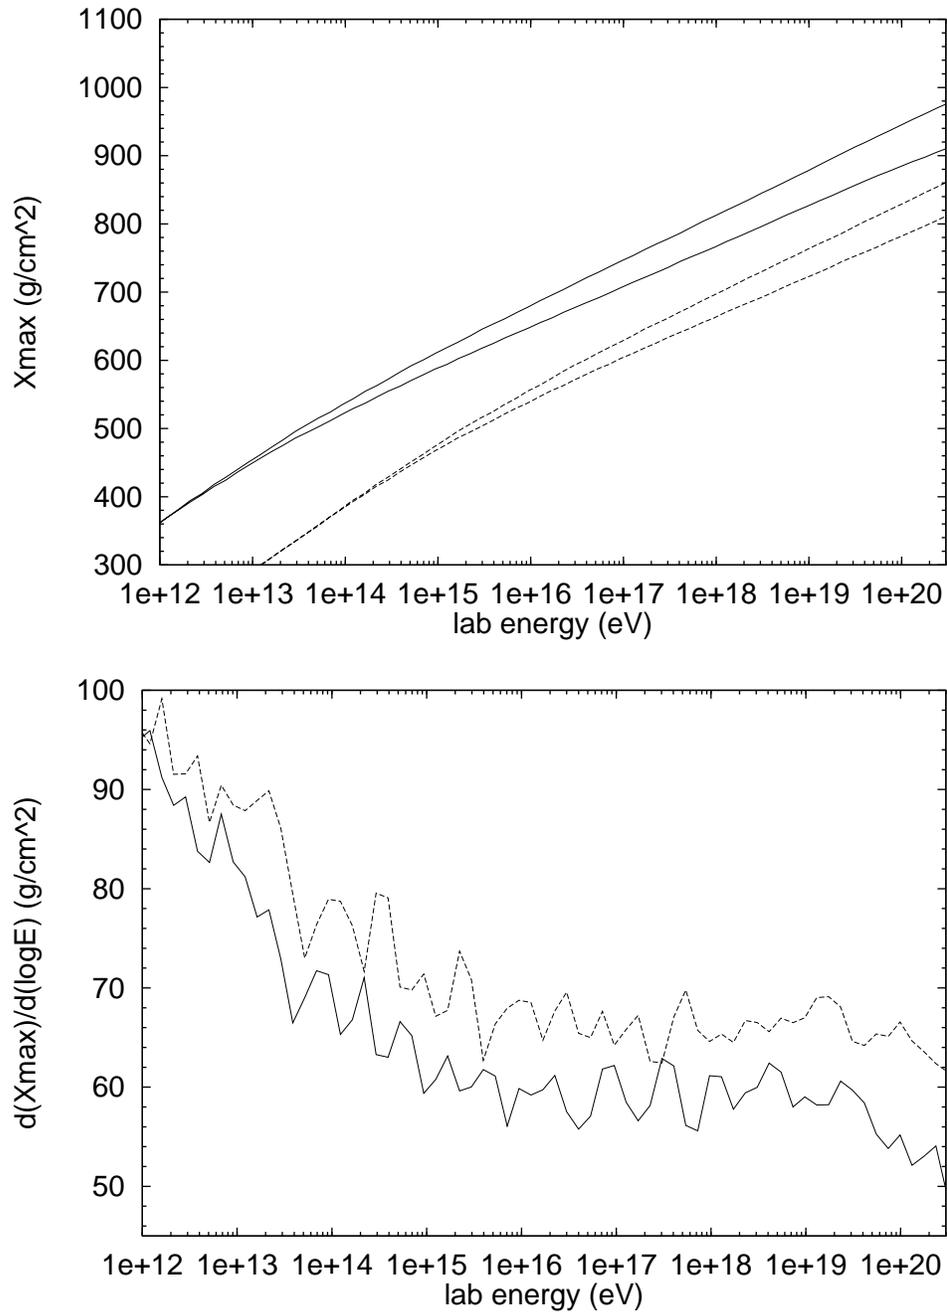


Figure 5.7: Using default cross sections, the energy dependence of the depth at maximum and the elongation rate for nucleons with only scaling violations in pp interactions turned on. The dashed lines in the depth of maximum plot are for iron. The results are compared with the baseline model (the upper line of each pair).

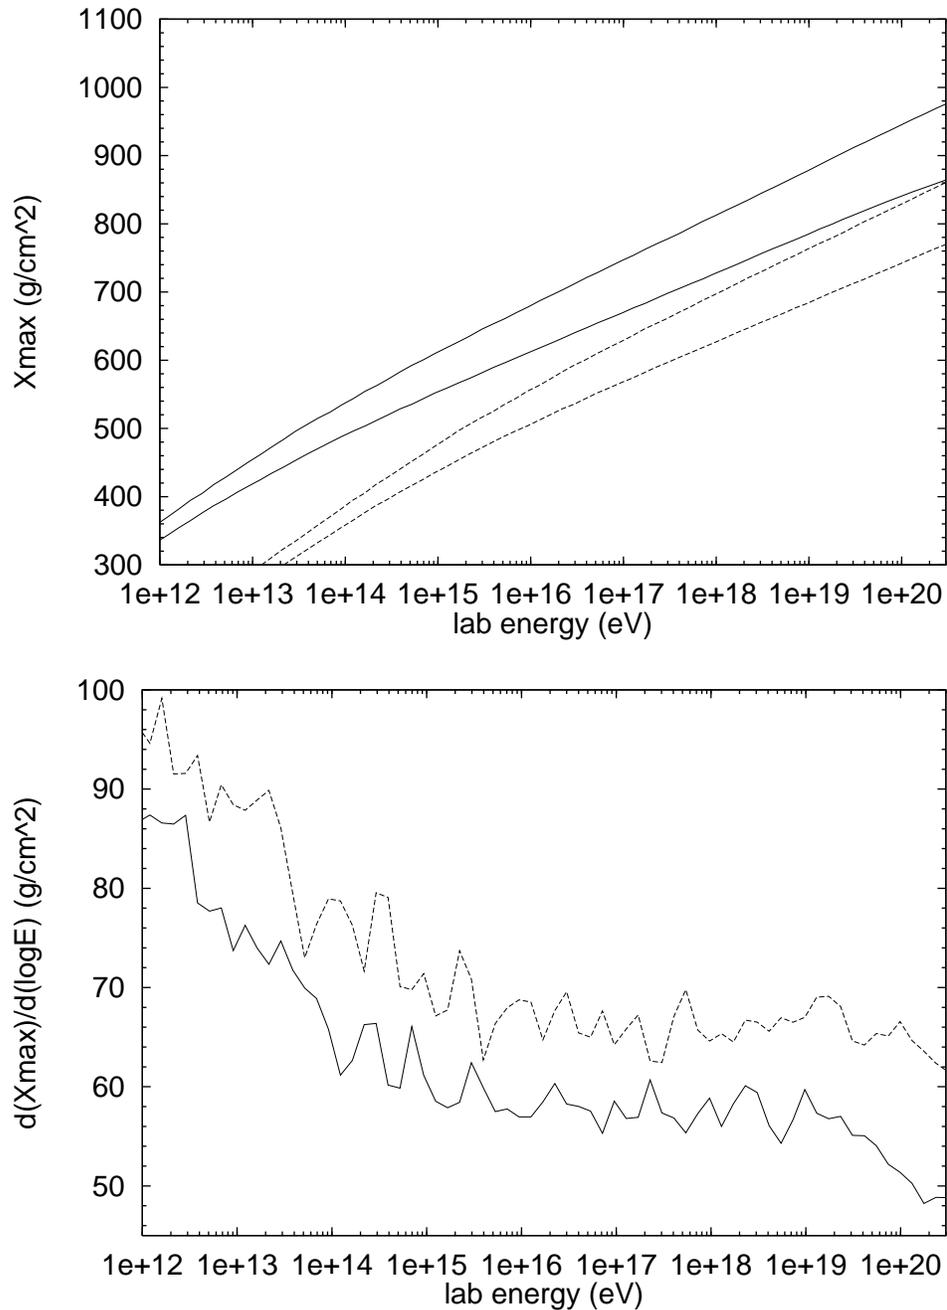


Figure 5.8: Using default cross sections, the energy dependence of the depth at maximum and the elongation rate for nucleons with scaling violations in pp interactions and nuclear target effects turned on. The dashed lines in the depth of maximum plot are for iron. The results are compared with the baseline model, the upper line of each pair.

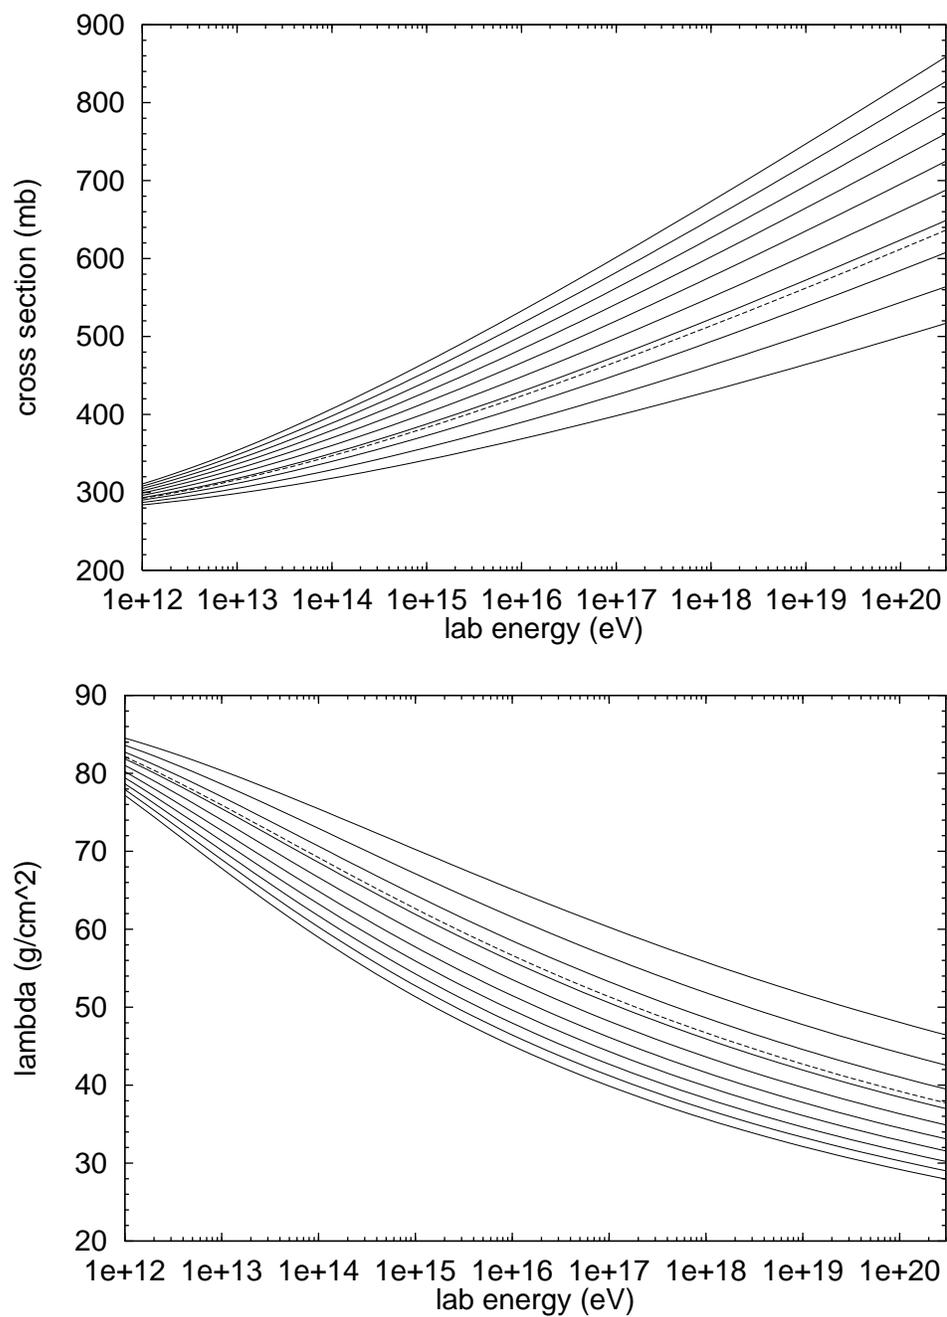


Figure 5.9: The energy dependence of the inelastic cross section (top) and the interaction length (bottom) for nucleons on air over the set of models investigated. The dotted line shows the result for default cross sections.

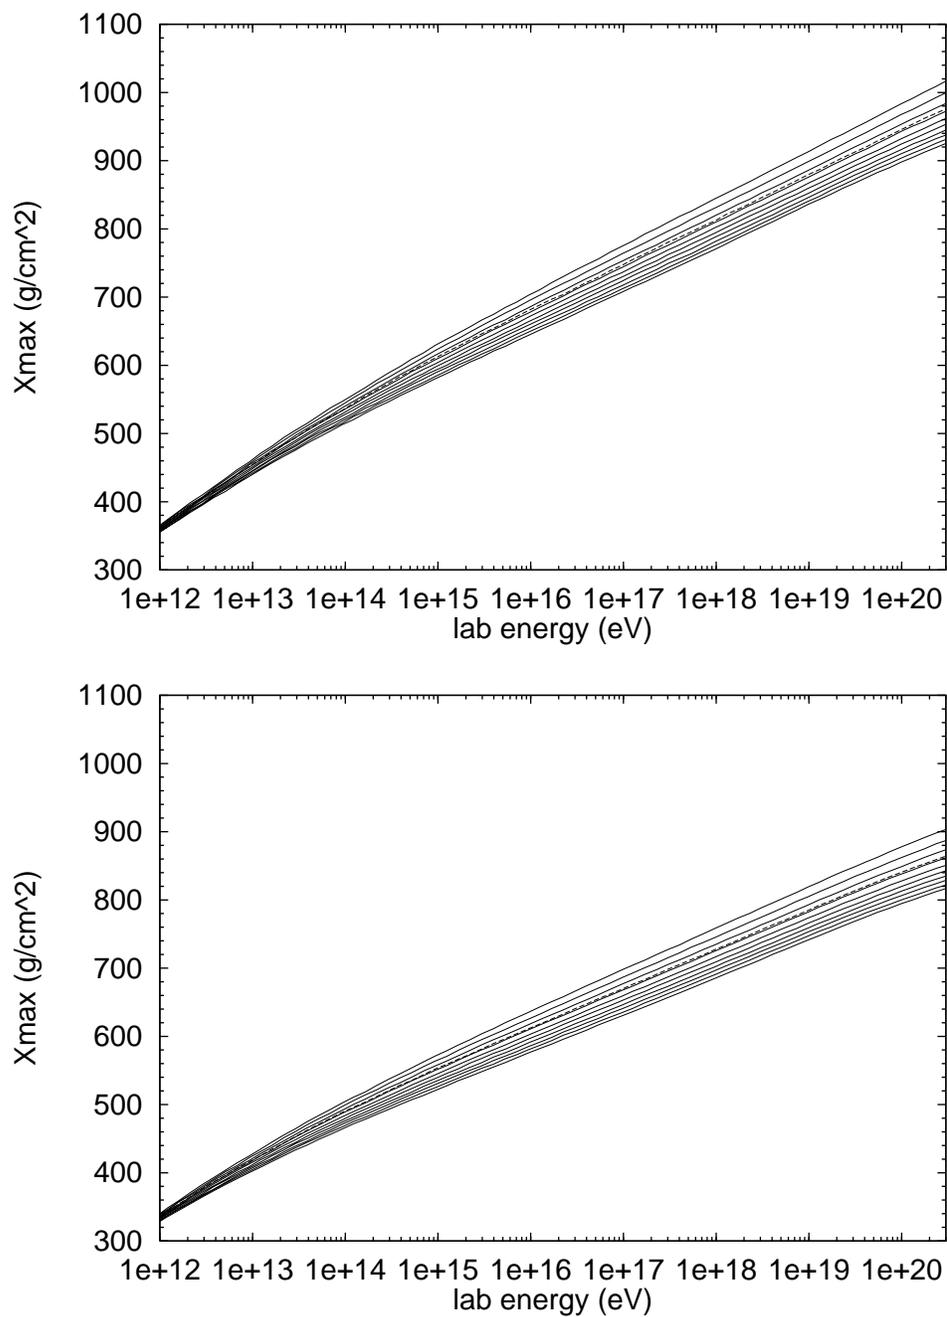


Figure 5.10: The energy dependence of the depth at maximum for the scaling model of hadronization (top graph) and with scaling violations in pp interactions and nuclear target effects turned on (bottom graph) over the set of models investigated. The dotted line shows the result for the default cross sections.

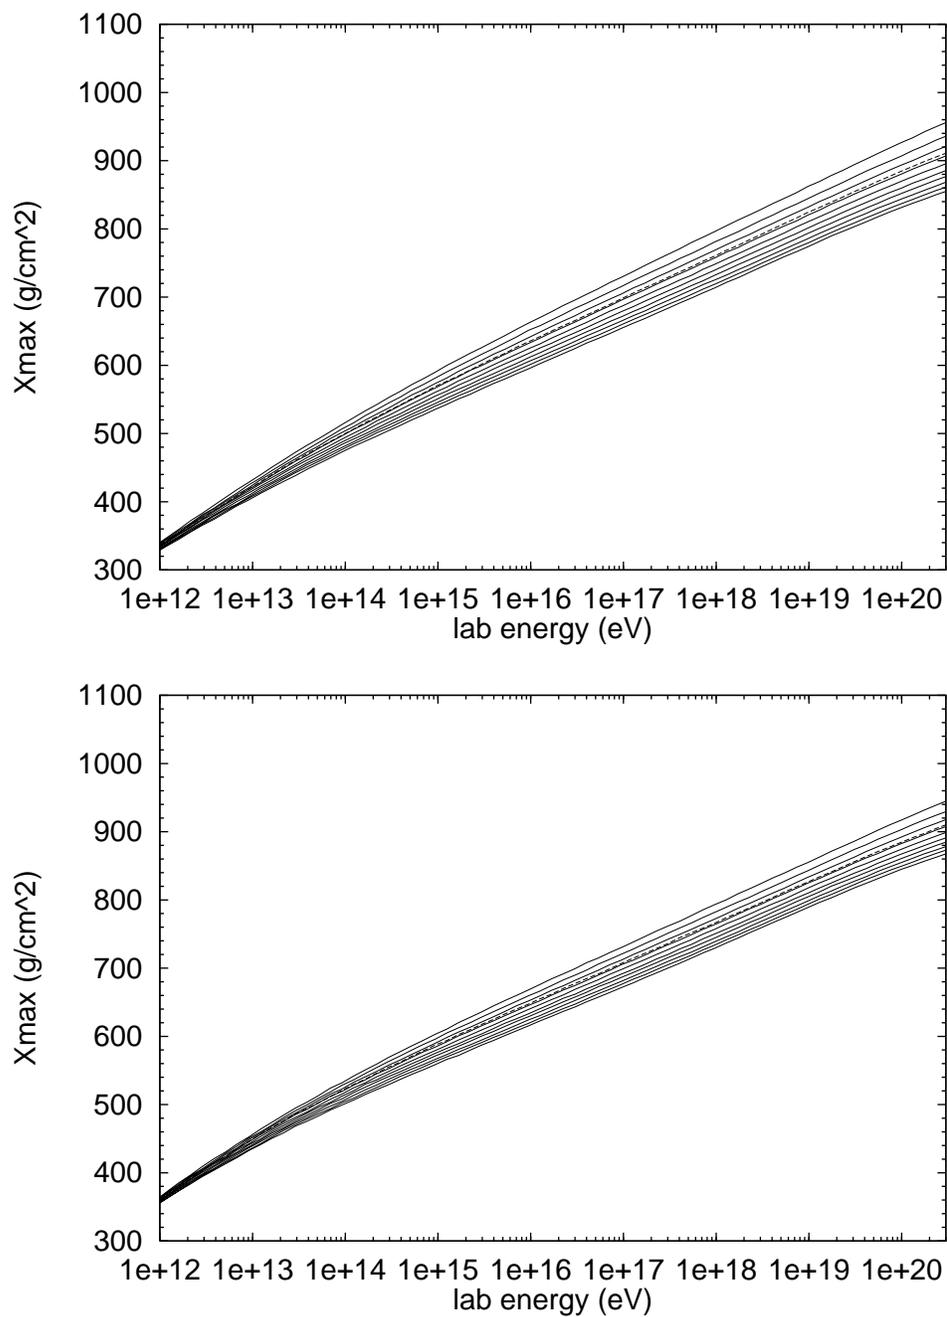


Figure 5.11: The energy dependence of the depth at maximum distribution with only nuclear target effects turned (top graph) and with only scaling violations in pp interactions turned on (bottom graph) over the set of models investigated. The dotted line shows the result for the default cross sections.

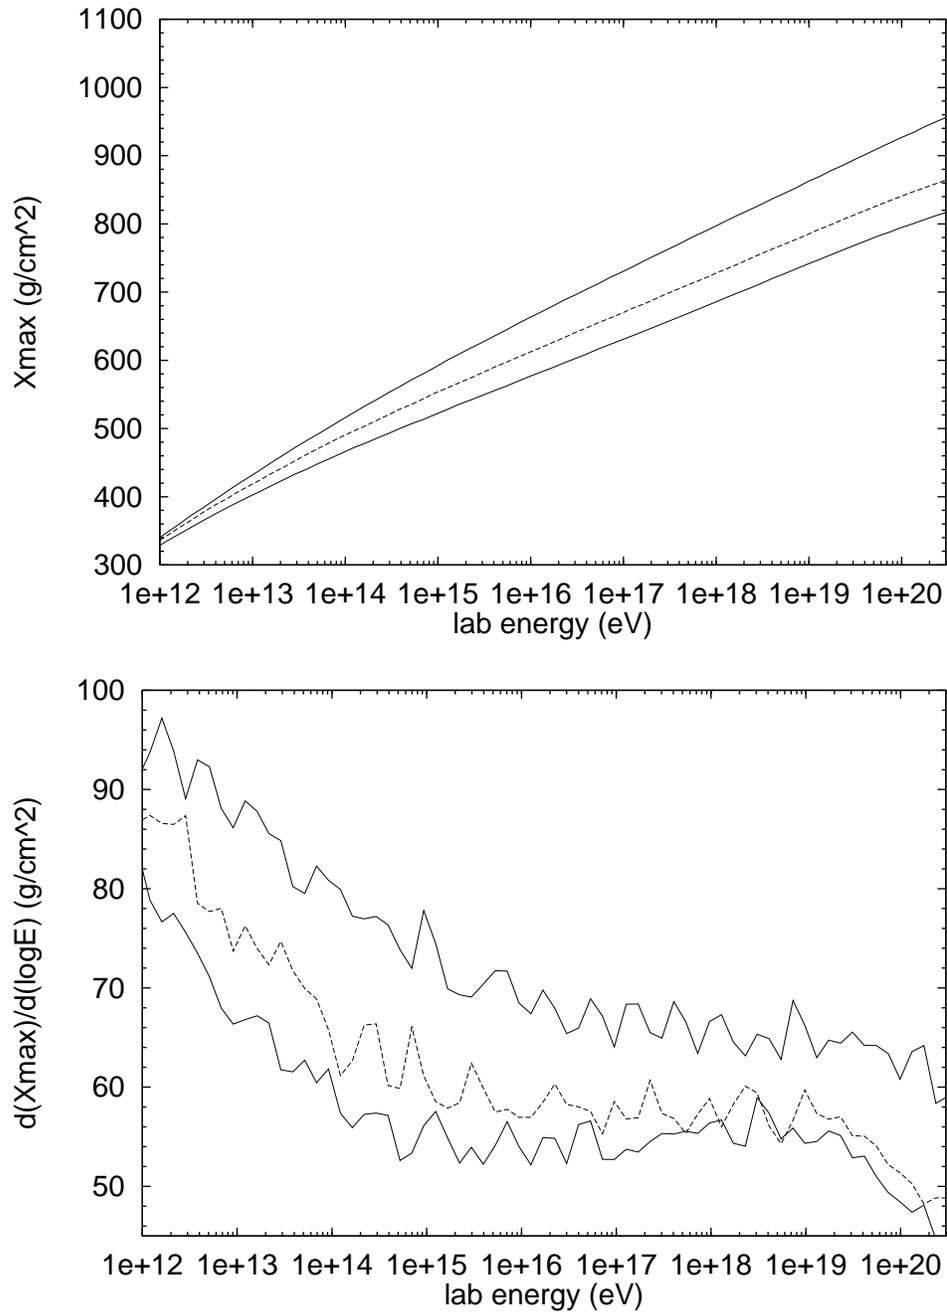


Figure 5.12: Bracketing the uncertainties in air shower development assuming proton primaries, energy dependence of the depth at maximum and the elongation rate. The dashed line is a lower bound from theory, the lower solid line is the lower bound constraint from air shower data.

Bibliography

- [1] “The Pierre Auger Project Design Report,” Fermilab, 1996 (unpublished).
- [2] J. F. Krizmanic and R. Streitmatter, “Signal, Background and Energy Threshold Estimates for the OWL Experiment”, Goddard Space Flight Center, OWL Note GSFC 6, (1997).
- [3] D. J. Bird et al., (Fly’s Eye Collab.) *Phys. Rev. Lett.* **71**, 3401 (1993; *Astrophys. J.* **424**, 491 (1994); *ibid.* **441**, 144 (1995).
- [4] S. Yoshida, et al., (AGASA Collab.) *Astropart. Phys.* **3**, 105 (1995); N. Hayashida et al., *Phys. Rev. Lett.* **73**, 3491 (1994).
- [5] G. Brooke et al. (Haverah Park Collab.), Proc. 19th Intl. Cosmic Ray Conf. (La Jolla) **2**, 150 (1985); reported in M. A. Lawrence, J. O. Reid, and A. A. Watson (Haverah Park Collab.), *J. Phys. G* **17**, 733 (1991).
- [6] N. N. Efimov et al., (Yatkutsk Collab.) ICRR Symposium on Astrophysical aspects of the Most Energetic Cosmic Rays, ed. N. Nagano and F. Takahara, World Scientific pu. (1991); and Proc. 22nd ICRC, Dublin (1991).
- [7] K. Greisen, *Phys. Rev. Lett.* **16**, 748 (1966).
- [8] G. T. Zatsepin and V. A. Kuzmin, *Sov. Phys. JEPT Lett.* **4**, 78 (1966).
- [9] S. Yoshida and M. Teshima, *Prog. Theor. Phys.* **89**, 833 (1993).
- [10] R. J. Protheroe and P. Johnson, *Astropart. Phys.* **4**, 253 (1996).

- [11] J. L. Puget, F. W. Stecker and J. H. Bredekamp, *Atrophys. J.* **205**, 638 (1976).
- [12] M. S. Longair, *High Energy Astrophysics Vols. 1-2*, Cambridge University Press, 2nd Edition (1992).
- [13] P. Sokolsky, *Introduction to Ultrahigh Energy Cosmic Ray Physics*, Addison-Wesley Publishing Company (1989).
- [14] T. K. Gaisser, *Cosmic Rays and Particle Physics*, Cambridge University Press (1992).
- [15] E. Fermi, *Phys. Rev.*, **75**, 1169 (1949).
- [16] M. Takeda, N. Hayashida, K. Honda et al., *Phys. Rev. Lett*, **81**, 1163 (1998).
- [17] G. Domokos and S. Kovesi-Domokos, *Phys. Rev. Lett*, **82**, 1366 (1999).
- [18] H. A. Bethe and W. Heitler, *Proc. Roy. Soc.*, **146**, 83 (1934).
- [19] W. Heitler, *Quantum Theory of Radiation*, Oxford University Press, 2nd edition (1944).
- [20] *Phys. Rev. D, Review of Particle Properties*, **50**, 1251 (1994).
- [21] B. Rossi, *High Energy Particles*, Prentice-Hall (1952).
- [22] A. M. Hillas, *J. Phys. G: Nucl. Phys.*, **8**, 1461 (1982).
- [23] T. K. Gaisser, *Workshop on Observing Giant Cosmic Ray Air Showers from $> 10^{20}$ eV from Space*, ed. J. F. Krizmanic, 297 (1997).
- [24] B. Rossi and K. Greisen, *Revs. Mod. Phys*, **13**, 240 (1941).
- [25] K. Greisen, *Prog. Cosmic Ray Physics*, **3**, 1.
- [26] E. J. Fenyves et al., *Phys. Rev. D*, **37**, 649 (1988).
- [27] T. K. Gaisser et al., *Phys. Rev. D*, **47**, 1919 (1993).

- [28] M. Honda et al., *Phys. Rev Letters*, **70**, 525 (1993).
- [29] R. M. Baltrusaitis et al., *Phys. Rev. Letters*, **52**, 1380 (1984).
- [30] B. Z. Kopeliovich, N. N. Nikolaev and I. K. Potashnikova, *Phys. Rev. D*, **39**, 769 (1989).
- [31] T. Wibig and D. Sobczynska, *J. Phys. G: Nucl. Part. Phys.*, **24**, 2037 (1998).
- [32] A. M. Hillas, *Proc. 17th Int. Cosmic Ray Conf.*, **8**, 193 (1981).
- [33] H. Bengtsson and T. Sjostrand, *Comput. Phys. Commun.* **46**, 43 (1987).
- [34] A. E. Brenner et al., *Phys. Rev. D*, **26**, 1497 (1982).
- [35] M. Aguilar-Benitez et al., *Z. Phys. C*, **50**, 405 (1991).
- [36] R. S. Fletcher, T. K. Gaisser, P. Lipari and T. Stanev, *Phys. Rev. D*, **50**, 5710 (1994).
- [37] F. Halzen and A. D. Martin, *Quarks & Leptons, An Introductory Course in Modern Particle Physics*, John Wiley & Sons (1984).
- [38] T. K. Gaisser and T. Stanev, *Phys. Lett.*, **219**, 375 (375).
- [39] C. Albajar et al., *Nucl. Phys.B*, **309**, 405.
- [40] W. M. Morse et al., *Phys. Rev. D*, **15**, 66 (1977).
- [41] G. J. Alner et al., *Phys. Lett.*, **167B**, 476 (1986).
- [42] R. J. Glauber and G. Matthiae, *Nucl. Phys. B*, **21**, 135 (1970).
- [43] J. Engel et al., —emph*Phys. Rev. D*, **46**, 5013 (1992).
- [44] L. K. Ding et al., *Astrophys. J.*, **474**, 490 (1997).
- [45] T. T. Chou and C. N. Yang, *Phys. Rev. D*, **32**, 1692 (1985).
- [46] M. Adamus et al., *Z. Phys. C*, **39**, 311 (1988).

- [47] G. D. Westfall et al., *Phys. Rev. C*, **19**, 1309 (1979).
- [48] R. Gandhi et al., *Astropart. Phys.*, **5**, 81 (1996).
- [49] T. K. Gaisser and A. M. Hillas, *Proc. 15th ICRC, Plovdiv, Bulgaria*, vol. **8**, 353 (1977).
- [50] G. M. Frichter, T. K. Gaisser and T. Stanev, *Phys. Rev. D*, **56**, 3135 (1997).
- [51] K. Werner, *Phys. Reports*, **232**, 87 (1993).
- [52] L. D. Landau and I.J. Pomeranchuk, *The Collected Papers of L. D. Landau*, Pergamon Press (1965).
- [53] A. B. Migdal, *Phys. Rev.*, **103**, 1811 (1956).
- [54] S. R. Klein, *Workshop on Observing Giant Cosmic Ray Air Showers from $> 10^{20}$ eV from Space*, ed. J. F. Krizmanic, 132 (1997).
- [55] N. N. Kalmykov, S. S. Ostapchenko and A. I. Pavlov, *Phys. Atomic Nuclei*, **58**, 1728 (1995).

Vita

Paul Mikulski was born in Baltimore, MD on April 21, 1969. He received his bachelor's degree from St. Mary's College of Maryland. In 1992 he came to Johns Hopkins University to pursue graduate studies under Professor Domokos and Professor Kovesi-Domokos. Beginning in 1996, he took a three-semester leave of absence to teach physics as a visiting instructor at St. Mary's College of Maryland. He is currently doing molecular dynamics research at the United States Naval Academy.